

AZURE HYBRID EXPERIENCE DAY 2024

Herzlich willkommen!

Agenda

09:00 Uhr: Ankommen & Registrierung

09:30 Uhr: Willkommen

Moderator: Daniel Schleidt, Leiter Wirtschaftsredaktion Rhein-Main, Frankfurter Allgemeine Zeitung

09:45 Uhr: Keynote: Adaptive Cloud verändert das Business - wie profitieren Sie davon?

Stefan Hörner, CEO, Medialine AG

10:15 Uhr: Azure Stack HCI & Azure Arc - alles auf einen Blick

Dino Bordonaro, Lead Architect Microsoft Multicloud Solutions, Microsoft MVP Azure, Medialine AG

11:00 Uhr: Pause & Raumwechsel

Thomas Maurer, Principal Program Manager, Azure Adaptive Cloud Experience, Microsoft



Agenda

11:30 Uhr: Azure Stack HCI & Azure Arc - alles auf einen Blick Part 2

Dino Bordonaro, Lead Architect Microsoft Multicloud Solutions, Microsoft MVP Azure, Medialine AG
Thomas Maurer, Principal Program Manager, Azure Adaptive Cloud Experience, Microsoft

12:00 Uhr: Einblicke in den Maschinenraum von Azure Stack HCI

Thomas Roettinger, Principal PM Manager, Microsoft

12:45 Uhr: Mittagspause

Track 1

Adaptive Cloud & Azure Stack HCI im Detail
Raum Redmond

Track 2

Adaptive Cloud & Azure Arc im Detail
Raum Chicago

Track 1
Raum Redmond



Manfred Helber

Hardwareauswahl - was
gilt zu beachten
Azure Stack Integration
ins eigene
Datacenter



Thomas Roettinger

Sicherheit in allen
Ebenen,
von der Hardware bis
zur App

Track 2
Raum Chicago

13:30 Uhr



Hardwareauswahl - was gilt zu beachten

Azure Stack Integration ins eigene Datacenter

Manfred Helber | Microsoft Most Valuable Professional

Azure Stack HCI Solutions



Premier Solutions

Turnkey Azure Stack HCI solution

- Deepest integration and highest level of automation, built through deep engineering collaboration between Microsoft and solution partners
- Continuous testing by Microsoft and our partners, to ensure higher reliability and minimal downtime
- End-to-end deployment workflows that make it easy to deploy one cluster or a thousand clusters

Integrated Systems

Single purpose system with pre-installed software

- Optimized hardware selection with regular testing for ongoing reliability
- Delivered with software pre-installed and security set by default
- Validated full-stack updates and native hardware management tools

Validated Nodes

Broadest choice of hardware components

- Choose from a diverse selection of validated hardware from more than 30 partners, or re-use existing validated hardware
- Engage with preferred SI for deployment and integration, as needed
- On new hardware - OR - Check with your OEM or solution provider to ensure you are running a validated solution. In certain cases, you may be able to re-use existing hardware

Azure Stack HCI Solutions


Feature	Validated Nodes	Integrated Systems	Premier Solutions
Validated solution with certified hardware configuration	✓	✓	✓
Solution committed to 5 years of hardware support	✓	✓	✓
Solution testing requirement	Once	2-4 times/year	Continuously
Ease of deployment and updates	Manual	Double-click	Single-click, seamless
HCI software pre-installed or on-site deployment services		✓	✓
Solution available as multi-node		✓	✓
Support from Microsoft and hardware solution partner		✓	✓
Security settings and features enabled by default		✓	✓
Validated by Microsoft in our own labs			✓
One stop for Level 1 and Level 2 support for HCI software			✓
Call home support service option available			✓
Global availability (solution and services in 100+ countries)			✓
White glove deployment services available			✓
As-a-service purchasing option available			✓
Optimal customer use case	Customers who run multiple operating systems or want to manage their own firmware, driver, and OS updates	Customers seeking some level integration and validation from Azure Stack HCI for multi-node clusters	Customers requiring turnkey, enterprise-wide deployments across their global portfolio with simplified management



Live Demo: Azure Stack HCI Catalog



Azure Stack HCI – Server-side architecture




Azure Stack HCI – Drive requirements



NVMe for Capacity

OR




NVMe for Cache SSD for Capacity

OR




SSD for Capacity

OR




NVMe for Cache HDD for Capacity

OR



SSD for Cache HDD for Capacity

OR




NVMe for Cache SSD + HDD for Capacity

Live Demo: Azure Stack HCI Sizer



4th and 5th Gen Intel® Xeon® processors for Microsoft




Shared **architectural platform** for 4th Gen Intel® Xeon processors and 5th Gen Intel® Xeon® processors

- ← Tremendous **performance** gains, particularly when modernizing from servers with older generation Intel® Xeon® processors
- ← Introduction of **built-in accelerators** on CPU to help maximize performance and power efficiency across a range of workloads
- Enhanced **performance per watt** while also enjoying improved performance
- Greater **support for AI** with Intel® AMX



Microsoft Azure Stack HCI Using 4th Gen Intel® Xeon® Scalable Processors Delivers Performance for Demanding Workloads

IOPS Improvement on Azure Stack HCI
VMfleet 2.0 (Higher is Better)



Performance increased gen-to-gen, scales with cluster size and higher series of Xeon



Common Training and Inference Workloads Accelerated by Intel® Advanced Matrix Extensions (Intel® AMX)

Image Classification

Speech Recognition


Language Translation

Object Detection



Why Run AI Workloads on Intel® Xeon® Processors with Microsoft Azure Stack HCI?

Any AI Code, Every Workload



A broad array of modern business applications are now being augmented with AI.

4th & 5th Gen Intel® Xeon® processors feature Intel® Advanced Matrix Extensions (Intel® AMX), which gives your AI-enabled apps the ability to deliver flexible and efficient performance.


Build & Deploy Everywhere



Design and deploy AI projects quickly and efficiently with optimized training and inferencing.

Intel AMX brings extensive hardware and software optimizations to enable fast and efficient AI for a range of use cases, including video analytics, industrial machine vision, and natural language processing.

Implement Pre-Built Solutions



With the built-in AI accelerator, the result is an optimized pipeline on a single hardware and software platform that scales from data center to cloud to edge.

Customers can scale AI everywhere by leveraging the broad, open, optimized libraries, frameworks and pre-trained models available to speed deployment.

The flexibility of a standard Intel® Xeon® Scalable processor server, with the efficiency and performance of a built-in AI accelerator


Accelerate AI - Image Classification on Microsoft Azure Stack HCI

using 4th Gen Intel® Xeon® Scalable Processors with

TensorFlow
Higher is better

Intel® AMX

Image Classification on Tensorflow 2.11 using ResNet50
BS=128, Multi-instance (16x2 and 40x2 instances)



- The ResNet-50 benchmark measures image classification/vision workloads
- FP32 is a standard 32-bit floating point data type used to train deep learning models and for inferencing
- Bfloat16 is a truncated version of 32-bit floating point, used for both training and inference, offering similar accuracy but faster computation
- INT8 offers higher performance and is least computationally demanding for constrained environments, with minimal impact on accuracy
- Many DL workloads are mixed precision and 4th Gen Intel Xeon Scalable processors can seamlessly transition between Intel AMX and Intel AVX-512 to use the most efficient instruction set


Increase performance with higher series of Intel Xeon processor or by changing precision

Accelerate Natural Language Processing (NLP) on Microsoft Azure Stack HCI using 4th Gen Intel® Xeon® Scalable Processors with Intel® AMX

 TensorFlow
Higher is better

NLP on Tensorflow 2.11 using BERT-Large

BS=128, Multi-instance (16x2 and 40x2 instances)




- BERT-Large is a pretrained model used for Natural Language Processing
- FP32 is a standard 32-bit floating point data type used to train deep learning models and for inferencing
- Bfloat16 is a truncated version of 32-bit floating point, used for both training and inference, offering similar accuracy but faster computation
- INT8 offers higher performance and is least computationally demanding for constrained environments, with minimal impact on accuracy
- Many DL workloads are mixed precision and 4th Gen Intel Xeon Scalable processors can seamlessly transition between Intel AMX and Intel AVX-512 to use the most efficient instruction set

Increase performance with higher series of Intel Xeon processor or by changing precision

Why Modernize Your Database Infrastructure with Intel® Xeon® Processors on Microsoft Azure Stack HCI?

Performance Efficiencies



Databases running on the latest Azure Stack HCI software with 4th and 5th Gen Intel® Xeon® Scalable processors can process more transactions with increased performance as compared to older clusters.

Better Decisions



Updating your Microsoft SQL database solution enables faster analytics on your wealth of data for near real-time decision making.

Data Protection



Modernizing your HCI and Microsoft SQL solution makes it easier to protect data by backing up more quickly, using fewer CPU resources, avoiding impact to database users.

When running SQL 2022, utilize built-in **Intel® QAT** to accelerate backup.

Upgrade to the latest software running on 4th Gen Intel® Xeon® Scalable processors for increased performance, faster analytics, while protecting valuable data.

Why Modernize Your VDI Infrastructure with Microsoft AVD on Azure Stack HCI with Intel® Xeon® Processors?

Address Today's Workforce Needs



Support your evolving, distributed, hybrid workforce with a modern VDI solution for low-latency, more secure computing from anywhere.

Reduce response times and speed data access, even with power and memory-hungry office, media, and web apps.

Simplified Management



Centralized management allows IT administrators to scale resources efficiently while staying on budget.

Increase flexibility to meet different user needs with Intel's scalable compute, GPU, networking, and storage portfolio.

Secure Access



Allow remote workers to securely access the tools they need, when and where they need them.

Enhance security by locking down data and apps in the datacenter vs. on vulnerable endpoints.

Microsoft Azure Virtual Desktop on Azure Stack HCI running on Intel® Xeon® processors delivers more user sessions, or more virtual CPU and memory per user, than previous generations

Game-changing Security from Intel and Microsoft

- ✓ Fortify your infrastructure
- ✓ Help protect data at Rest, in Transit, and in Use
- ✓ Encrypt everything

Secure hardware, firmware and OS capabilities to help protect against threats

Intel Supports Secured Core

Reduces the performance impact of pervasive encryption, making it possible for users to broadly encrypt data with minimal performance impact

Intel® Crypto Acceleration Technology

Encrypts memory holistically across the platform using hardware-generated keys, and encrypts VMs to help protect existing software


Intel® Total Memory Encryption and Total Memory Encryption-Multi-Key

Accelerating Cryptography and Data De/compression – supported in Microsoft SQL 2022 on Azure Stack HCI



Intel® QuickAssist Technology





Azure Stack HCI Networking




Azure Stack HCI Networking




Azure Stack HCI Networking




Klassische TCP/IP Kommunikation (ohne RDMA)




Traditional mode



Netzwerk-Kommunikation mit RDMA




RDMA mode



RDMA Options

- RoCE:
 - RoCE is based on the IB proven technology
 - Lighter weight protocol than iWARP (L2 vs. L3)
 - Lowest latency
 - Requires loss-less Ethernet (DCB / PFC)
 - Scaling concerns over congestion-spreading issues
 - Non-routable, no IP header
- RoCE v2:
 - Adds an IP/UDP header for Routing
 - Improved Network management (for IP networks)
 - Still requires loss-less Ethernet (DCB / PFC)
- iWARP:
 - Employs TCP/IP as its transport
 - Does not require loss-less Ethernet (DCB / PFC) to function
 - Somewhat higher latency, but scales to Internet proportions

Partners with certified RDMA adapter




Partners with certified RDMA adapter



Broadcom NetXtreme P2200G
10/25/50/100/200G



Intel E810 10/25/50/100G



Nvidia (Mellanox) ConnectX-7
10/25/50/100/200/400G



Marvell (Qlogic/Cavium) FastLinQ QL45611HLCU
10/25/50/100G

Die aktuell flexibelste NIC für ASHCI



Intel E810 10/25/50/100G

Intel® Ethernet 800 Series – What's New

10/25/100Gb Controllers and Adapters

Key New Features

Precision Timing

- High-accuracy timing synchronization to ensure continuous services and maintain revenue streams
- SyncE Delivers frequency synchronization over Ethernet between sync enabled devices. Extends traditional telecom timing architecture to commercial-off-the-shelf (COTS) Ethernet solutions

Open vSwitch (OVS) Acceleration

- DPDK integration with OVS increases performance by eliminating extra layers in the OVS architecture/native stack

RDMA RoCEv2 LAG

- Link aggregation increase bandwidth
- Provides high-availability and resiliency
- Load balancing distributes traffic across physical links, provides resiliency and network redundancy



25Gb Intel® Ethernet Network Adapter E810-XXVDA4T w/GNSS



100Gb Intel® Ethernet Network Adapter E810-CQDA2T w/GNSS

Hyperconverged Infrastructure (HCI)

Azure Stack HCI

Support both iWARP and RoCEv2 RDMA


3.06M IOPS
>2GB/Sec D Transfer
0.4MS Response Time

Demonstrates support for the most demanding business needs: streaming, file access, LTP and other use cases. With RDMA, the performance was sustained even as the cluster size increased.

For more information go to intel.com/ethernet



Networking und Firewall Requirements für ASHCI



Raumwechsel

14:15 Uhr

Track 1
Raum Redmond



Marcel Meurer

Azure Virtual Desktop -
von Azure ins
eigene RZ



Jan Philipp Hoepfner

AKS Hybrid -
Azure Kubernetes im
eigenen RZ

Track 2
Raum Chicago

14:20 Uhr



Netzwerk Pause

Kaffee &
Snacks



Track 1
Raum Redmond



Peter Sommer

Migration, Backup und
Disaster Recovery
Möglichkeiten auf einen
Blick



Manfred Helber

Betrieb- und Self-Service
von Azure
Stack HCI powered by
Azure Arc

Track 2
Raum Chicago


15:45 Uhr




Betrieb- und Self-Service von Azure Stack HCI powerd by Azure Arc

Manfred Helber | Microsoft Most Valuable Professional

Azure Stack HCI



Azure Stack HCI – Deployment von Workloads



Live Demo: Betrieb- und Self-Service von Azure Stack HCI powered by Azure Arc



Raumwechsel

16:30 Uhr

A person wearing a white long-sleeved shirt stands against a clear blue sky. Instead of a head, there is a large, fluffy white cloud. The person's arms are slightly raised, and their hands are open.

16:35 Uhr

**Closing Keynote und Roadmap
Adaptive Cloud: Copilot for Azure &
Arc & HCI, VisionAi & IoT-Ops**

Martin Hörhammer, Dino Bordonaro, Medialine AG

Thomas Maurer, Sven Langenfeld, Microsoft





17:30 Uhr

Get together: Beer, Wine & Good Talks

Zur Umfrage

Ende!

Vielen Dank für Ihre Teilnahme!

Medialine EuroTrade AG

sales@medialine.ag
www.medialine.com

Breitlerstraße 43
55566 Bad Sobernheim
Tel: +49 6751 85378-0
Fax: +49 6751 85378-11

www.group.medialine.com



MEDIALINE
GROUP