Beyond Hyperconverged: New Azure Local Capabilities on Lenovo ThinkAgile MX

Solution Brief

Organizations want the agility of the cloud without giving up control over where their data lives and how their infrastructure performs. Microsoft Azure Local answers that need with Azure-consistent management and services running on infrastructure the customer owns.

Recent Azure Local releases have significantly widened what that infrastructure can do. New options for enterprise storage integration, GPU-accelerated workloads, physical resiliency, simplified identity, and confidential computing extend Lenovo ThinkAgile MX beyond a traditional hyperconverged platform. Depending on configuration, ThinkAgile MX can now serve as the compute foundation for disaggregated designs, host AI and visualization workloads, span two racks or rooms as a single cluster, and isolate regulated workloads with hardware-based memory encryption.

This brief describes those capabilities, the business needs they address, and the ThinkAgile MX systems that support them.

Business Challenge

Hybrid infrastructure teams face a set of pressures that a conventional hyperconverged design does not always resolve:

- **Uneven growth**: Compute and storage demand rarely grow at the same rate. In a hyperconverged cluster, adding capacity of one kind means buying the other, and existing SAN investments sit outside the Azure-connected estate.
- **AI close to the data**: Inference and visualization workloads need GPUs near the point where data is created, yet GPU resources are often managed separately from the rest of the environment and cannot move during maintenance.
- **Rack-level failures**: A single power, switching, or rack fault can take down an entire cluster. Running two clusters to avoid that doubles the management burden.
- **Edge overhead**: Stores, branches, and plants frequently have no Active Directory infrastructure, and adding domain controllers at every site increases cost, complexity, and attack surface.
- **Data in use**: Regulated industries already encrypt data at rest and in transit, but sensitive workloads remain exposed in memory to host administrators and compromised host software.

Each of these gaps slows hybrid cloud adoption or pushes workloads onto infrastructure that is harder to govern.

This document is for employee-use only

Solution

Lenovo ThinkAgile MX combines validated Lenovo ThinkSystem server technology with Azure Local in rack configurations (ThinkAgile MX630 V4, MX650 V4, MX650a V4, MX630 V3, and MX650 V3) and a compact edge system (ThinkAgile MX455 V3 Edge). ThinkAgile MX Premier Solutions are co-engineered and continuously tested by Lenovo and Microsoft, with one Lenovo contact for hardware and software support. Validated firmware and drivers are delivered through the Lenovo Solution Builder Extension, lifecycle is coordinated through Lenovo XClarity and Azure tools, and Lenovo Open Cloud Automation scales deployment from one site to a large fleet.

Five capabilities through Azure Local 2608

- **External SAN storage (GA)**: ThinkAgile MX V4 compute connects to ThinkSystem DM, DS, or DG Series arrays over Fibre Channel or iSCSI with Multipath I/O. Hybrid HCI plus SAN, fully disaggregated, and existing SAN models are supported, so compute and storage scale independently.
- **GPU provisioning and mobility (GA)**: Dedicated GPUs (DDA) or GPU partitions (GPU-P) are assigned during or after VM creation, and GPU-P workloads live migrate between compatible nodes. GPU management from Azure is in preview.
- **Rack-aware clustering (GA)**: One Azure Local cluster spans two racks or rooms, each a local availability zone, with synchronous storage replication between them.
- **Local Identity with Azure Key Vault (GA)**: Nodes deploy without joining an Active Directory domain, using certificate-based authentication, while credentials and BitLocker recovery keys are protected in Azure Key Vault.
- **Confidential VMs (preview)**: On the ThinkAgile MX455 V3 Edge with AMD EPYC processors, AMD SEV-SNP encrypts VM memory and isolates it from the host, with Microsoft Azure Attestation gating the release of secrets.

ThinkAgile MX is available as a traditional purchase or as a service through Lenovo TruScale, backed by Lenovo Professional Services, Premier Support, and Managed Services.

This document is for employee-use only

Uses Cases

The capabilities described above matter most where a conventional hyperconverged design has forced a compromise: buying storage to get compute, building domain infrastructure at every site, or accepting a single rack as a point of failure. The scenarios below show how organizations in different industries apply them on ThinkAgile MX, often combining more than one capability in the same deployment.

These scenarios are illustrative rather than exhaustive. Your Lenovo representative or Lenovo Business Partner can help map them to your own sites, workloads, and growth plans.

AI inference at retail and branch locations

Retailers and distributed service organizations want computer vision, recommendation, and analytics models running in the store rather than in a distant region. A compact ThinkAgile MX deployment with GPU acceleration runs inference where data is created, and Local Identity with Azure Key Vault removes the need to stand up Active Directory at every site. Central IT manages GPUs and clusters from the same Azure experience used across the estate.

Manufacturing and operational technology

Plant floor systems must keep running, and OT networks are often deliberately isolated from corporate domains. ThinkAgile MX deploys with Local Identity, so no domain infrastructure is required on site. Rack-aware clustering places nodes in two separate racks or rooms, so production workloads continue through the loss of power, switching, or an entire rack, without the overhead of operating two clusters.

Regulated workloads in healthcare, finance, and public sector

Organizations bound by privacy and sovereignty rules need sensitive data to stay on-premises and stay protected while it is processed. Confidential VMs on ThinkAgile MX455 V3 Edge isolate workload memory from host administrators, and hardware attestation verifies the environment before keys are released. External SAN storage adds the enterprise data services these organizations already rely on.

Datacenter modernization with existing SAN

Enterprises with established Fibre Channel or iSCSI storage can bring that capacity into Azure Local rather than retire it. ThinkAgile MX V4 compute nodes attach to qualified ThinkSystem DM, DS, or DG Series arrays, existing storage practices remain in place, and compute is added only when compute is needed.

Workloads

Not every capability is available on every ThinkAgile MX system, and several arrived in later Azure Local releases rather than the initial ones. Planning a deployment therefore starts with matching each capability to a supported ThinkAgile MX model, the minimum Azure Local release, and the configuration rules that Microsoft and Lenovo have validated together.

Lenovo delivers the corresponding firmware, drivers, and utilities through the Lenovo Solution Builder Extension, aligned with Lenovo Best Recipe releases, so that each node runs a tested combination. The summary below captures the key requirements for each capability:

- **External SAN (Fibre Channel from 2604, iSCSI from 2607)**: ThinkAgile MX630 V4, MX650 V4, and MX650a V4. Requires redundant fabrics or dedicated Ethernet paths, MPIO, identical adapters across nodes, and NTFS Cluster Shared Volumes. Disaggregated SAN-only instances scale to 64 machines.
- **GPU acceleration (DDA and GPU-P, from 2604)**: ThinkAgile MX455 V3 Edge, MX650 V3 Premier Solution with controlled GPU, MX650 V4, and MX650a V4. GPU-P live migration requires homogeneous GPUs and NVIDIA vGPU software 18 or later.

This document is for employee-use only

- **Rack-aware clustering (from 2601)**: 1+1, 2+2, 3+3, or 4+4 nodes across two zones; all-flash storage; dedicated storage network of 10 GbE minimum (25 GbE recommended); 1 ms or less round-trip latency between racks; new deployments only.
- **Local Identity with Azure Key Vault (GA from 2604)**: All supported ThinkAgile MX solutions; static IP addressing, DNS, and one Key Vault per instance. Windows Admin Center is not supported in this model.
- **Confidential VMs (preview from 2607)**: ThinkAgile MX455 V3 Edge Premier Solution with 4th Gen AMD EPYC processors and Ubuntu 24.04 guests.

The following figure shows a 2+2 rack-aware cluster on ThinkAgile MX: each room has its own pair of top-of-rack switches, and storage traffic is replicated synchronously between rooms so that either rack can carry the workload if the other fails.

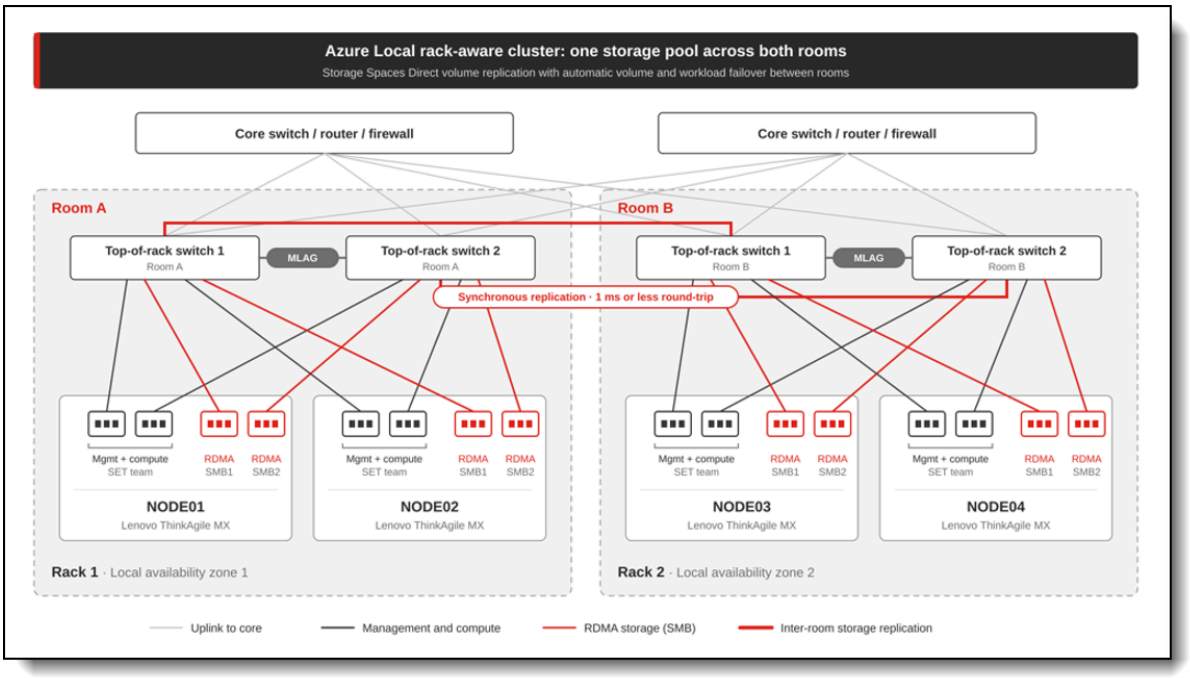

Figure 1. Azure Local rack-aware cluster network architecture (adapted from Microsoft Learn)

This document is for employee-use only

Business Outcome

For IT leaders, the value of these capabilities lies less in any single feature than in how far they extend one consistent platform. Workloads that previously required separate infrastructure, such as SAN-backed applications, GPU-accelerated AI, edge sites without Active Directory, and regulated data, can now run on ThinkAgile MX under the same Azure-connected operating model, with one Lenovo contact for support.

That consolidation means fewer platforms to purchase, staff, patch, and secure, and a shorter path from a new business requirement to a running workload. The outcomes break down as follows:

- **Pay for what grows**: Disaggregated designs scale compute and storage independently, up to 64 machines per instance compared with 16 for a hyperconverged instance, and put existing SAN capacity back to work.
- **Higher availability**: Rack-aware clusters tolerate the failure of one complete rack or zone while keeping a single deployment to manage.
- **Faster edge rollout**: Removing Active Directory as a prerequisite cuts the infrastructure each site needs, and Lenovo Open Cloud Automation repeats deployments across a fleet.
- **AI without new silos**: GPU workloads are provisioned and maintained through the same Azure tools as every other VM.
- **Lower risk**: Co-engineered Premier Solutions, validated SBE packages, and a single point of support shorten time to value.

Performance and TCO figures for these configurations have not yet been published; your Lenovo representative can help size and model a deployment for your environment.

Conclusion

Azure Local releases through 2608 extend Lenovo ThinkAgile MX from hyperconverged infrastructure into disaggregated storage, GPU-accelerated AI, rack-level resiliency, AD-free edge deployment, and confidential computing. Because support depends on the exact ThinkAgile MX model and Azure Local release, Lenovo and its Business Partners help map each capability to the right configuration, and Lenovo Professional Services can plan and deliver the deployment, whether purchased outright or consumed through Lenovo TruScale.

This document is for employee-use only

For More Information

To learn more about Azure Local capabilities on Lenovo ThinkAgile MX, contact your Lenovo representative or Lenovo Business Partner, or visit the resources below.

References:

- Lenovo Press: Implementing Azure Local Capabilities Through Release 2608 on Lenovo ThinkAgile MX (Implementation Guide): https://lenovopress.lenovo.com/lp2567-azure-local-capabilities-on-thinkagile-mx
- Lenovo Press: Lenovo ThinkAgile MX650 V4 Hyperconverged System Product Guide: https://lenovopress.lenovo.com/lp2137-lenovo-thinkagile-mx650-v4-hyperconverged-system
- Lenovo Press: ThinkSystem DM7200F SAN with ThinkAgile MX650 V4 in Azure Local: https://lenovopress.lenovo.com/lp2385-solution-brief-dm7200f-san-with-thinkagile-mx650-v4-in-azure-local
- Lenovo Press: Lenovo Configuration Guide for Confidential Computing and Microsoft CVM: https://lenovopress.lenovo.com/lp2484-lenovo-configuration-guide-for-confidential-computing-and-microsoft-cvm
- Lenovo Docs: ThinkAgile MX Solution Builder Extension overview: https://pubs.lenovo.com/thinkagile-mx/mx_sbe_overview
- Microsoft Learn: What's new in Azure Local: https://learn.microsoft.com/en-us/azure/azure-local/whats-new?view=azloc-2608
- Microsoft Learn: Azure Local disaggregated deployments overview: https://learn.microsoft.com/en-us/azure/azure-local/overview/disaggregated-overview

Authors

Chris Honoré is a Solutions Product Manager at Lenovo with deep expertise in datacenter products and solution offerings. He has a strong background in consulting and solution development, helping customers design and support on-premises and hybrid environments. Chris has spent the past 15 years with IBM and Lenovo, specializing in x86 server and data center solutions. Prior to that, he built two decades of experience in the telecommunications industry, serving in both technical and business leadership roles.

Adelina Onofrei is a Microsoft Solutions Engineer, specializing in Azure Local deployments, Azure Local cluster and LXCI testing. Adelina is dedicated to enabling business success through secure, innovative, and reliable solutions based on Azure Local and enterprise cloud technologies.

Related product families

Product families related to this document are the following:

- DG Series Storage
- DM Series Storage
- DS Series Storage
- Intel Alliance
- Microsoft Alliance
- ThinkAgile MX Series for Microsoft Azure Local

This document is for employee-use only

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION ”AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2026. All rights reserved.

This document, LP2579, was created or updated on November 22, 2026.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
  https://lenovopress.lenovo.com/LP2579
- Send your comments in an e-mail to:
  comments@lenovopress.com

This document is available online at https://lenovopress.lenovo.com/LP2579.

This document is for employee-use only

Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at https://www.lenovo.com/us/en/legal/copytrade/.

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
Lenovo TruScale®
ThinkAgile®
ThinkSystem®
TruScale®
XClarity®

The following terms are trademarks of other companies:

AMD®, AMD EPYC®, and EPYC® are trademarks of Advanced Micro Devices, Inc.

Microsoft®, Active Directory®, Azure®, BitLocker®, NTFS®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

NVIDIA® is a trademark of NVIDIA Corporation.

Other company, product, or service names may be trademarks or service marks of others.

This document is for employee-use only