



Running SQL Server with High Availability on ThinkAgile MX with Azure Stack HCI Stretched Cluster Feature

Last update: **06 April 2021**

Version 1.0

Highlights benefits of Azure Stack HCI stretched cluster for SQL Server high availability configurations

Presents a use case for Lenovo ThinkAgile MX offerings

Includes deployment information and best practices

Contains detailed bill of materials for servers

**Laurentiu Petre
David West
Vinay Kulkarni**



Table of Contents

1	Introduction	1
2	Business value	2
3	Architectural overview Stretched Clusters	4
4	Prerequisites for Azure Stack HCI Stretched cluster	6
5	Cluster creation and configuration with Windows Admin Center	7
5.1	Cluster creation.....	7
5.2	Cluster configuration	7
6	Architectural overview SQL Server high availability	8
7	Prerequisites for SQL Server high availability	9
8	Appendix: Bill of Materials	10
8.1	Server BOM	10
	Resources	13

1 Introduction

This document describes the benefits of running Microsoft SQL Server workloads on an Azure Stack HCI stretched cluster solution from Lenovo. Microsoft released the first version of Azure Stack HCI operating system name 20H2 in December 2020. One of shiny new features in Azure Stack HCI is stretched cluster. Azure Stack HCI stretched clustering on Lenovo ThinkAgile MX series provides you a disaster recovery solution for your mission critical SQL Server workloads.

The intended audience is IT professionals, technical architects, sales engineers, and consultants to assist in planning, designing, and implementing this solution.

This document provides an overview of the business problem and business value that is addressed by having a highly available SQL instance on an Azure Stack HCI stretched cluster.

2 Business value

High availability and disaster recovery are important to keep SQL Server workloads and data safe and available online during planned and unplanned downtimes. A good disaster recovery solution can cost a lot of money to implement but can help an organization maintain their customers trust and earn significant revenue. A company that is planning to implement a disaster recovery solution can reduce its costs by adding an instance of SQL on the Azure stack HCI stretch cluster without any performance penalty if sized correctly.

Lenovo ThinkAgile MX offering for Azure Stack HCI provides enterprise customers a highly available, cost efficient, flexible platform to run a high-performance Microsoft SQL Server leveraging the power of state-of-the-art hardware and Storage Spaces Direct. Azure Stack HCI presents a highly competitive solution for delivering exceptionally performant Microsoft SQL Server. Whether running Online Transaction Processing (OLTP) workloads, or Data Warehouse and BI, to AI and advanced analytics over Big Data, you will benefit from the resiliency that Azure Stack HCI offers. This is especially important for mission critical databases. Leveraging the flexibility to run SQL Server in VMs (Windows Server or Linux), it allows you to consolidate multiple database workloads and easily scale out by adding additional VMs to the Azure Stack HCI environment as needed.

Additionally, Azure Stack HCI enables you to integrate Microsoft SQL Server with Azure Site Recovery service to provide a cloud-based migration, restoration, and protection solution for your organization's data that is reliable and secure.

Lenovo has multiple solutions that are ideal for consolidation of legacy SQL Server deployments on ThinkSystem servers. Lenovo offers these solutions under the ThinkAgile offerings as ThinkAgile MX for Azure Stack HCI.

[ThinkAgile MX](#)

Lenovo rack systems feature innovative hardware, software and services that solve customer challenges today and deliver an evolutionary fit-for-purpose, modular design approach to address tomorrow's challenges. These servers capitalize on best-in-class, industry-standard technologies coupled with differentiated Lenovo innovations to provide the greatest possible flexibility in x86 servers. Key advantages of deploying Lenovo rack servers include:

- Highly scalable, modular designs to grow with your business
- Industry-leading resilience to save hours of costly unscheduled downtime
- Expansive storage capacity and flexible storage configurations for optimized workloads

Fast flash technologies for lower latencies, quicker response times and smarter data management in real-time for cloud deployments, database, or virtualization workloads, trust Lenovo racks for world-class performance, power-efficient designs, and extensive standard features at an affordable price.

The following Lenovo servers have been certified for Microsoft Azure Stack HCI and are equipped to support up to 2x 8-64-core processors, up to 4TB of memory and over 100TB of storage making them ideal candidates for Azure Stack HCI SQL Server solutions:

- Lenovo ThinkAgile MX3520 Integrated systems and ThinkAgile MX validated nodes (based on ThinkSystem SR650)
- Lenovo ThinkSystem SR665 validated nodes
- Lenovo ThinkSystem SR630 validated nodes

3 Architectural overview Stretched Clusters

Microsoft Azure Stack HCI Stretched cluster is a robust solution for setting up a disaster recovery with automatic failover of the resources. Storage replica allows synchronization between volumes in different sites

Figure 1 shows an architectural overview of a stretched cluster with two nodes in each site. It is recommended to have at least two nodes in each site for intra-site failover.

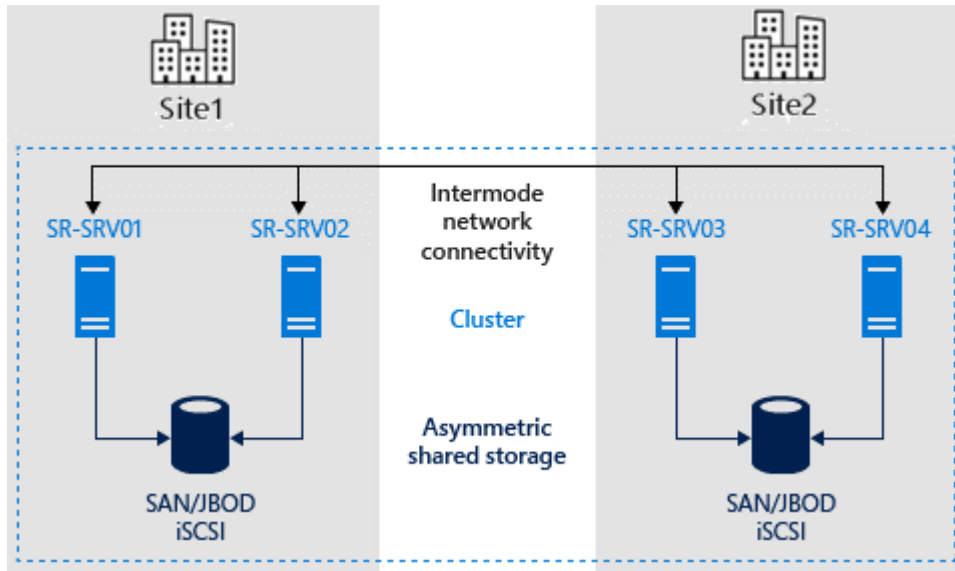


Figure 1: Typical stretched cluster with two nodes in each site

The stretched clusters can be divided in two categories based on the replication type: *active-passive* and *active-active*. In a *active-active* case the replication will be performed both ways between the sites and in a *active-passive* the replication will be performed in one direction only.

In figure 2 we have the active-passive diagram and in figure 3 we have the active-active diagram.

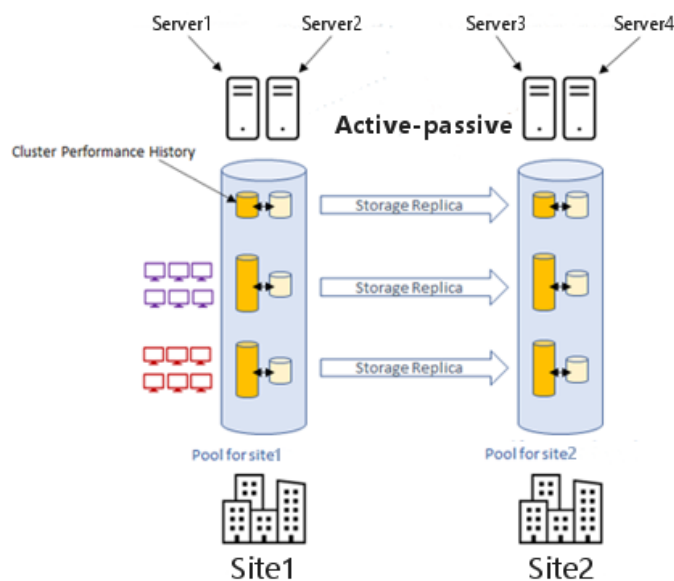


Figure 2: Active-passive diagram

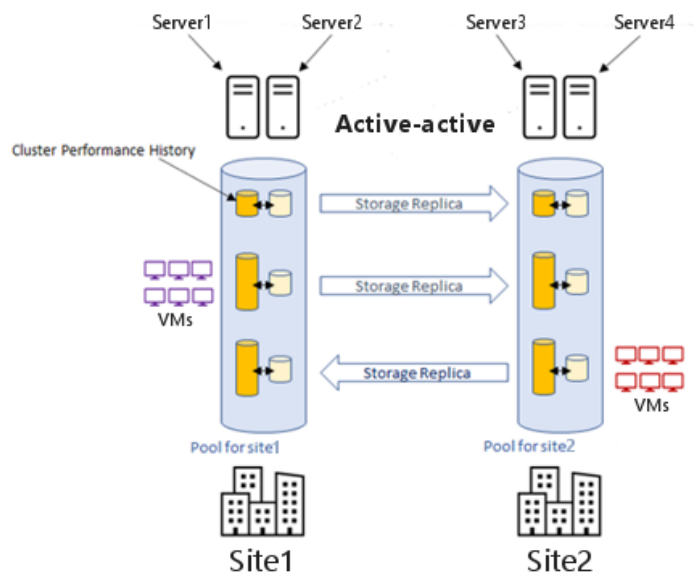


Figure 3: Active-active diagram

More, in depth information, can be found in Microsoft's [documentation here](#).

4 Prerequisites for Azure Stack HCI Stretched cluster

A typical Azure Stack HCI cluster requires a minimum of two server and in the case of a stretch cluster it is recommended to have two nodes per site in order to achieve intra-site failover.

The servers will need to be deployed in two separate sites configured in Active Directory. There are no physical distance restrictions regarding the sites if the connectivity and latency is ensured.

As in the case of any cluster, it is recommended that all the servers to have the same hardware configuration on all the nodes and that the RAM and storage is sized accordingly to the workloads that will be configured.

The cluster can be created by using Windows Admin Center that is installed on a PC or a Server.

5 Cluster creation and configuration with Windows Admin Center

5.1 Cluster creation

The Azure Stack HCI Stretched Cluster can be configured from Windows Admin Center by following the dedicated wizard and making sure to select the *server in two sites* option at the beginning. The detailed, step by step process is presented in [this article](#) created by Microsoft

5.2 Cluster configuration

In this scenario we will use two volumes (one in each site) that will not be set to replicate between sites. On these volumes the SQL virtual machines will be located, and the high availability will be achieved at the application layer. Therefore, when creating the volumes for SQL select the default “*create in a single site*” option.

6 Architectural overview SQL Server high availability

This section describes the mapping of the logical components of Microsoft SQL Server onto Lenovo servers, storage, and networking. The BOM configurations for the hardware are described in the Appendix.

The solution uses SQL Server's Availability Group (AG) feature to manage data replication at the SQL layer to provide database level high availability and disaster recovery. SQL is installed on the VMs, which are clustered at the VM level. Although not covered in this document, it is also possible to use SQL Failover Cluster Instances (FCI) on VMs to provide SQL application level protection along with AGs.

SQL AG's support synchronous or asynchronous replication. Synchronous is slightly slower but is always in synch with no data loss at failover. Asynchronous is faster but results in some data loss during failover.

The SQL Enterprise edition supports up to 8 replicas (copies) per group, and an unlimited number of availability groups.

Additional readable copies can be configured for access by read-only workloads, reporting, or to perform backups from them.

Failover can be performed manually or configured to trigger automatically if the primary database becomes unavailable. After a failover occurs, the secondary copy becomes the primary copy and clients are automatically redirected to the new copy as primary.

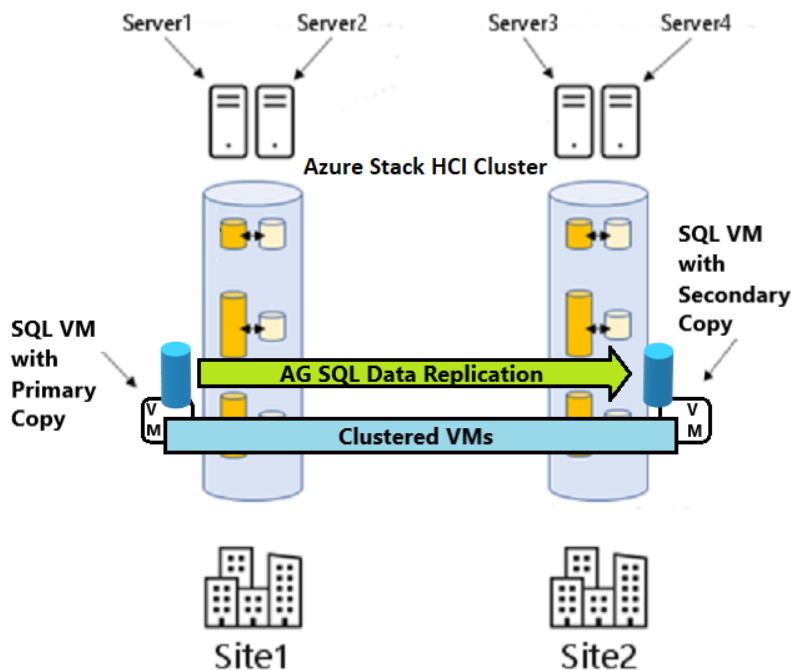


Figure 4 SQL Availability Group on VMs overview

7 Prerequisites for SQL Server high availability

This section contains additional configurations and prerequisites for deploying Microsoft SQL Server on a stretch cluster.

A key requirement is that each SQL instance run on a member of a Windows Server Failover Cluster, as SQL Availability Groups rely on Windows clustering. In this case, the VMs are guest clusters.

After the Windows cluster is formed with a VM from each site, an instance of SQL is installed on each VM.

Availability Groups are configured using SQL Server Management Studio. This includes creation of an AG listener object in Active Directory and DNS to direct clients to the current primary copy of the database.

Clients and applications should connect to the database using the AG Listener IP or its DNS name, to ensure the client always reconnects to the primary database if failover occurs.

The SQL data paths must be identical on each server, the AG setup process verifies this during setup. Below is a look at the SQL Server Management Studio where AG's are configured.

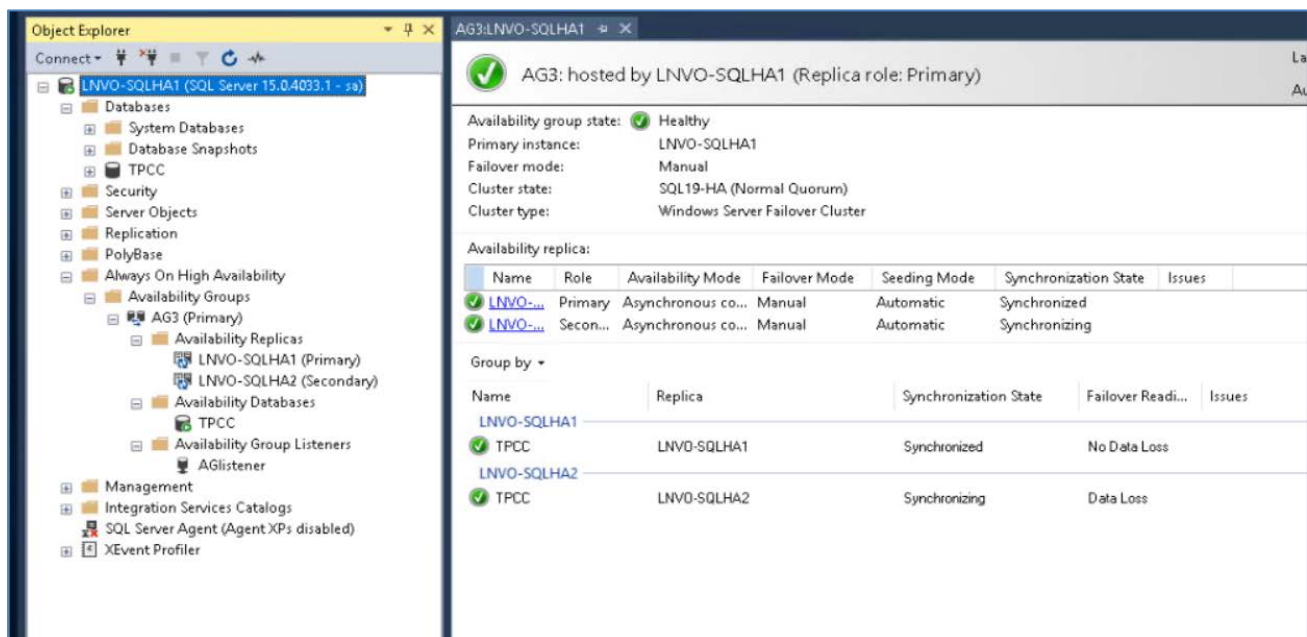


Figure 5 View of SQL SSMS with configured Availability Group and databases

Additional detailed information on prerequisites and configuration from Microsoft can be found [here](#)

8 Appendix: Bill of Materials

This appendix features the Bill of Materials (BOMs) for different configurations of hardware for Azure Stack HCI and Microsoft SQL Server deployments on Lenovo hardware.

The BOM listed in this appendix is not meant to be exhaustive and must always be confirmed with the configuration tools.

8.1 Server BOM

The following table list the server BOMs for a typical SQL Server workload (Needs to be adjusted based on customer requirements).

Table 1: ThinkAgile MX3520-F (All flash)

Part number	Product Description	Qty
7D5RCTO2WW	Server : ThinkAgile MX3520-F MX All-flash Appliance	1
B4E3	ThinkAgile MX Certified Node - All Flash	1
B4HT	Intel Xeon Silver 4208 8C 85W 2.1GHz Processor	2
B4H1	ThinkSystem 8GB TruDDR4 2933MHz (1Rx8 1.2V) RDIMM	12
B49B	ThinkSystem 2.5" Intel S4510 1.92TB Entry SATA 6Gb Hot Swap SSD	4
AURC	ThinkSystem SR550/SR590/SR650 x16/x8(or x16) PCIe FH Riser 2 Kit	1
AUR3	ThinkSystem SR550/SR650 x16/x8 PCIe FH Riser 1 Kit	1
AUAJ	Mellanox ConnectX-4 Lx 10/25GbE SFP28 2-port PCIe Ethernet Adapter	1
B919	ThinkSystem M.2 5300 480GB SATA 6Gbps Non-Hot Swap SSD	2
AUPW	ThinkSystem XClarity Controller Standard to Enterprise Upgrade	1
AVWF	ThinkSystem 1100W (230V/115V) Platinum Hot-Swap Power Supply	2
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
AXCA	ThinkSystem Toolless Slide Rail	1
B4NL	ThinkSystem SR650 Refresh MB	1
9220	Preload by Hardware Feature Specify	1
AURA	ThinkSystem 2U/Twr 2.5" SATA/SAS 8-Bay Backplane	2
AUR5	ThinkSystem 2U/Twr 2.5" AnyBay 8-Bay Backplane	1
5977	Select Storage devices - no configured RAID required	1
AUNL	ThinkSystem 430-8i SAS/SATA 12Gb HBA	3
B7XZ	Disable IPMI-over-LAN	1
AUMV	ThinkSystem M.2 with Mirroring Enablement Kit	1
BFXP	SW stack MS Azure Stack HCI	1
AUK7	TPM 2.0 and Secure Boot	1
AURD	ThinkSystem 2U left EIA Latch Standard	1
AUTJ	ThinkSystem common Intel Label	1
AUSA	Lenovo ThinkSystem M3.5" Screw for EIA	4
AVJ2	ThinkSystem 4R CPU HS Clip	2

Part number	Product Description	Qty
AURP	Lenovo ThinkSystem 2U 2FH Riser Bracket	1
AUSU	ThinkSystem Package for SR650	1
B51W	MX Series 2U Agency Label 2U	1
B13M	ThinkAgile EIA Plate	1
BFXQ	MX Appliance Badge	1
B0ML	Feature Enable TPM on MB	1
AWFF	ThinkSystem SR650 WW Lenovo LPK	1
AUAK	2U Bracket for Mellanox ConnectX-4 Lx 2x25GbE SFP28 Adapter	1
B7E5	MR customer specific - ROW	1
BD2F	MX Custom Config	1
AUSQ	On Board to 2U 8x2.5" HDD BP NVME Cable	1
AUSH	MS First 2U 8x2.5" HDD BP Cable Kit	1
AUSM	MS 2nd 2U 8X2.5" Cable Kit	2
AUSG	ThinkSystem SR650 6038 Fan module	1
AUTQ	ThinkSystem small Lenovo Label for 24x2.5"/12x3.5"/10x2.5"	1
AURR	ThinkSystem M3.5 Screw for Riser 2x2pcs and Planar 5pcs	4
AURY	Lenovo ThinkSystem PHY Module Dummy	1
AUT8	ThinkSystem 1100W RDN PSU Caution Label	1
B31F	ThinkSystem M.2 480GB SSD Thermal Kit	1
AVEP	ThinkSystem 4x1 2.5" HDD Filler	2
AVEN	ThinkSystem 1x1 2.5" HDD Filler	10
AURS	Lenovo ThinkSystem Memory Dummy	12
AURQ	Lenovo ThinkSystem 2U 3FH Riser Bracket	1
B173	Companion Part for XClarity Controller Standard to Enterprise Upgrade in Factory	1
AUTA	XCC Network Access Label	1
B13Q	ThinkAgile 2U Service Label LI	1
AWF9	ThinkSystem Response time Service Label LI	1
AUSE	Lenovo ThinkSystem 2U CPU Entry Heatsink	2
AUTU	ThinkSystem 4-7 NVMe sequence Label for 16x2.5"and 24x2.5"	1
AUTY	ThinkSystem 12-15 sequence Label for 24x2.5"Chassis	1
A2HP	Configuration ID 01	3
5374CM1	Configuration Instruction	1
AVE7	ThinkSystem 430-8i SAS/SATA 12Gb HBA placement	1
A2JX	Controller 01	1
A2HP	Configuration ID 01	1
5374CM1	Configuration Instruction	1
AVE7	ThinkSystem 430-8i SAS/SATA 12Gb HBA placement	1
A2JZ	Controller 03	1
A2HP	Configuration ID 01	1

Part number	Product Description	Qty
5374CM1	Configuration Instruction	1
AVE7	ThinkSystem 430-8i SAS/SATA 12Gb HBA placement	1
A2JY	Controller 02	1
A2HP	Configuration ID 01	1
5PS7A89731	Essential ThinkAgile APP - 3Yr 24x7 4Hr Resp + YDYD MX3520	1
	Auto-Derived Part Items	
B589	ThinkSystem U.2 Intel P4610 1.6TB Mainstream NVMe PCIe3.0 x4 Hot Swap SSD	2

Conclusion

Lenovo ThinkAgile MX is an ideal platform for configuring Microsoft SQL Server to run with high availability in a cost-effective model using the new stretched cluster feature only available in Microsoft Azure Stack HCI operating system.

Resources

For more information about the topics that are described in this document, see the following resources:

- Lenovo ThinkAgile MX for Azure Stack HCI.
[ThinkAgile MX](#)
- Microsoft SQL Server:
<https://docs.microsoft.com/en-us/sql/index>
- Microsoft Azure Stack HCI:
<https://docs.microsoft.com/en-us/azure-stack/hci/>

Trademarks and special notices

© Copyright Lenovo 2021.

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo, the Lenovo logo, ThinkSystem, ThinkAgile, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

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

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.