

Windows® Compute Cluster Server 2003: Top Linpack Benchmark Results

Mitsubishi UFJ Securities Cluster Achieves 6.5 TFLOPS Microsoft® Rainier Cluster Achieves 9.0 TFLOPS

Introduction

Windows® Compute Cluster Server (CCS) 2003 provides a full clustering ecosystem. Microsoft and its customers are building supercomputing clusters that scale with some of the largest and most powerful supercomputers in the world while providing the deployment, security, and development environments of Microsoft® Windows platform.

Mitsubishi UFJ Securities Cluster: 6.52 TFLOPS

Mitsubishi UFJ Securities Group is part of Mitsubishi UFJ Financial Group - one of the largest financial services institutions in Japan. Their expanding derivatives business will leverage the power of Windows based clustering to enhance risk management practices and reduce simulation times. Mitsubishi UFJ Securities chose the Microsoft Windows platform because of the power, familiarity, and ease-of-development in Microsoft Visual Studio® 2005 and Visual C++. Their Linpack benchmark on CCS 2003 with Service Pack 1 was run on a 448-node IBM BladeCenter HS21 cluster with 1,792 processors. The Linpack result of 6.52 TFLOPS demonstrates the power

of the advanced simulation environment that will enable the pricing, risk management, and product development of their derivatives offering.

Microsoft Rainier Cluster: 8.99 TFLOPS

The Microsoft Rainier cluster is a scalable test bed that Microsoft built in early 2007 to test high-performance software products, including the next version of CCS. Built with Dell PowerEdge 1955 blade servers, the Rainier cluster achieved a Linpack benchmark result of 8.997 TFLOPS on 256 nodes and 2,048 processor cores. The Microsoft Datacenter in Tukwila, Washington, which houses the Rainier cluster, gives Microsoft the ability to test software, hardware, and systems against a real-world, scalable cluster.

Linpack Benchmark

The Linpack benchmark is used by the Top500 project to measure the “best”

performance. The choice of Linpack as the benchmark for the Top500 supercomputers reflects its wide use and general availability on most relevant systems.

The Top500 uses a version of Linpack that allows the user to optimize the software by scaling the size of the problem to the cluster. As with any benchmark, Linpack is not an indicator of system performance under all conditions. Linpack has proven to reflect the system's performance in solving a dense system of linear equations.

Top500

The Top500 project (<http://www.top500.org>), which started in 1993, provides consistent and reliable tracking of high-performance computing (HPC) trends. Twice each year, the project assembles and releases a list of the 500 most powerful supercomputers in the world, as measured by the Linpack benchmark.

Rainier Cluster Configuration

The Rainier cluster used at the Microsoft Datacenter in Tukwila, Washington, consists of 26 Dell BladeCenter Chassis, with 10 Dell 1955 blades in each chassis for a total of 260 nodes. The networking used 22 Cisco SFS7000P SDR InfiniBand leaf-node switches, 12 Cisco SFS7000P SDR InfiniBand core switches, 260 Topspin InfiniBand HCA daughter cards, 26 InfiniBand I/O pass-through (one per chassis) for the InfiniBand networking, and a Cisco Catalyst 6509-E chassis populated with (7) Cisco Catalyst 6748 48-port line modules for Gigabit Ethernet (GigE). Cisco SFS7000P switches were configured in a core-leaf configuration. For the Linpack benchmark,

“I am pleased to see these new Top500 systems based on Windows Compute Cluster Server. The submissions underscore Microsoft's profound interest in our community and prove that supercomputing can be done on Windows.”

**Prof. Dr. Hans Werner Meuer,
Co-Founder of the Top500 list**

Mitsubishi UFJ Securities Cluster Achieves 6.5 TFLOPS Microsoft® Rainier Cluster Achieves 9.0 TFLOPS

a Microsoft Office Excel® 2007 add-in was used, along with the Microsoft Visual Studio 2005 Tools for the Microsoft Office System. The Office Excel 2007 add-in was used to submit and track jobs as well as tune and optimize the Linpack problem size for the Rainier cluster. This custom Office Excel 2007 add-in will be available shortly on the HPC Community site (<http://windowshpc.net/default.aspx>).

Rainier Cluster Deployment

The tools built into CCS 2003 provide an easy and efficient deployment mechanism with a familiar Windows interface. The head and imaging nodes were initially deployed with Windows Server 2003 with SP2, and Windows Server 2003 Automated Deployment Services (ADS) version 1.1 was installed

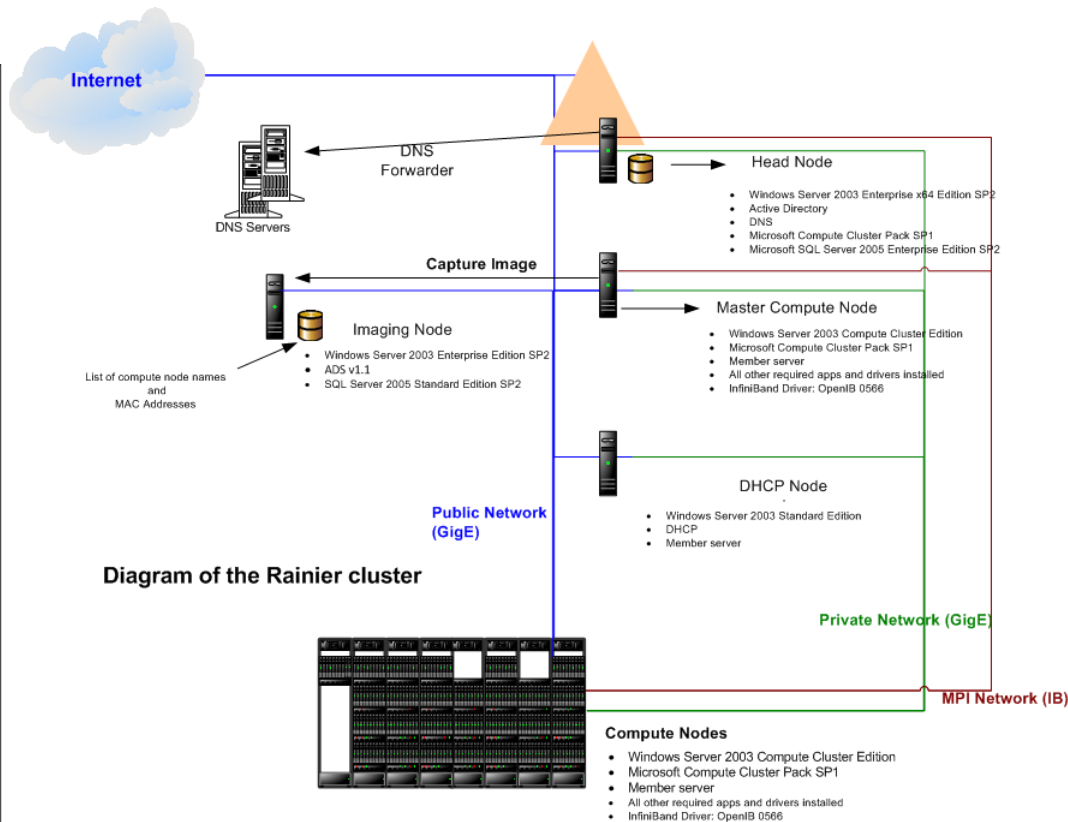
on the imaging node. The master compute node was created from CD, and the image was then deployed using multi-cast ADS to the other compute nodes. CCS 2003 was used to manage the cluster, schedule the compute jobs, and execute Message Passing Interface (MPI)–based applications, such as the Linpack benchmark.

“Microsoft has been making impressive inroads in HPC. Dell is pleased to be partnering with Microsoft to provide solid performance platforms for clustered environments. Dell’s HPC solutions combined with CCS can provide customers with an architectural model that supports growing HPC workloads and delivers the same ease of implementation and usage expected in a Windows-based environment.”

**Reza Rooholamini, Director,
Dell Software Solutions Engineering**

“IBM’s long standing high performance computing leadership on TOP500 is further enhanced by opening up the power of clusters to a whole new set of clients with the support of Microsoft Windows Compute Cluster 2003.”

IBM



The CCS 2003 Offer

CCS 2003 accelerates customers’ time-to-insight and productivity by providing a reliable HPC platform that is easy to deploy, operate, and integrate with existing infrastructure and tools. These Linpack benchmark results demonstrate that CCS scales from the workgroup cluster configuration all the way to some of the largest supercomputing configurations in the world.

For More Information

The CCS 2003 and HPC home page: <http://www.microsoft.com/hpc>

Windows HPC Community home: <http://windowshpc.net/default.aspx>

The Top500 home page: <http://www.top500.org>

Visual Studio 2005 Tools for the Microsoft Office System download: <http://www.microsoft.com/downloads/details.aspx?displaylang=en&FamilyID=f5539a90-dc41-4792-8ef8-f4de62ff1e81>

Porting instructions for Linpack: http://windowshpc.net/files/4/porting_unix_code/default.aspx

Mitsubishi UFJ Securities Cluster Achieves 6.5 TFLOPS Microsoft® Rainier Cluster Achieves 9.0 TFLOPS

Hardware	Microsoft Rainier	Mitsubishi UFJ Securities
Blades and chassis	(260) Dell PowerEdge 1955 blades (256 compute nodes); 26 PowerEdge chassis (10 blades per chassis)	IBM BladeCenter HS21
CPU	(2) x 1.86 GHz Intel Quad-Core 5320 Processor, (2) x 4 MB cache	(2) x 3.0 GHz Intel Dual-Core 5160 Processor, 4 MB cache
Bus	1,066 MHz FSB	1,333 MHz FSB
Chipset	Intel 5000P chipset	Intel 5000P chipset
RAM	(8) x 1 GB 667 MHz Dual-ranked DIMMS	(4) x 1 GB PC2-5300 Single-ranked FBDIMM
Local disk storage	SAS controller, (1) 73 GB SAS 2.5-inch 10K RPM disk	SAS controller 2.5-inch, 10K RPM disk
Gigabit Ethernet hardware	(1) Intel Gigabit Ethernet daughter card per blade, (2) Gigabit pass-through modules per chassis	(2) Broadcom Gigabit Ethernet 5708
	Cisco Catalyst 6509-E chassis with (7) 48-port Cisco 6748 line modules	Nortel Layer 2/3 Copper GigE Switch Module for IBM Blade-Center
	SFS7000P switches configured in a core-leaf configuration	
InfiniBand hardware	(22) Cisco SFS7000P SDR InfiniBand leaf-node switches	N/A
	(1) Topspin InfiniBand HCA daughter card per blade, (1) InfiniBand I/O pass-through for PowerEdge 1955 module per chassis	N/A
	(12) Cisco SFS7000P SDR InfiniBand core switches	N/A
Software		
Head node	Windows Server 2003 Enterprise x64 Edition with SP2 Microsoft SQL Server™ 2005 Enterprise Edition with SP2 Windows Compute Cluster Pack with SP1	Windows Server 2003 Standard x64 Edition with SP2 SQL Server 2005 Standard Edition with SP2 Compute Cluster Pack with SP1
Imaging node	Windows Server 2003 Enterprise Edition ADS v1.1 SQL Server 2005 Standard Edition	Windows Server 2003 Standard Edition R2 Altiris Deployment Solution 6.5 for IBM System X
Compute nodes	Windows Server 2003, Compute Cluster Edition, with SP2 Compute Cluster Pack with SP1 InfiniBand Driver: OpenIB 0566	Windows Server 2003, Compute Cluster Edition, with SP2 Compute Cluster Pack with SP1

© 2007 Microsoft Corporation. All rights reserved. This data sheet is for informational purposes only. No part of this document may be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without the express written permission of Microsoft Corporation. MICROSOFT MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS SUMMARY. Microsoft®, Excel®, Visual Studio®, Windows®, and the Windows logo are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. Other product and company names herein may be the trademarks of their respective owners.

Microsoft Corporation • One Microsoft Way • Redmond, WA 98052-6399 • USA

Microsoft®