



INSPIRE

breakthrough
availability with the
power of convergence

HP Enterprise Data Warehouse Appliance Optimized for SQL Server 2008 R2 Parallel Data Warehouse

Solution brief

The HP Enterprise Data Warehouse Appliance is an integrated solution optimized for SQL Server 2008 R2 Parallel Data Warehouse and designed for enterprises requiring large-scale data warehouse technologies in a pre-configured single enclosure. It incorporates the Microsoft application platform and HP Converged Infrastructure to provide organizations with

- Powerful data warehouse (DW) infrastructure delivering high performance, throughput, and capacity on an optimized set of server nodes and storage arrays.
- Greater scalability and massively parallel processing to support high-volume data demands driven by a hub-and-spoke architecture tuned specifically for distributed query processing.
- Reduced acquisition costs and faster time-to-deployment through an appliance-like form factor that stands-up in days rather than months.
- Easier ongoing operation enabled by end-to-end integration, intelligent control, and mission-critical support.
- Configuration built on industry-leading, standards-based HP hardware to integrate with your existing data center and meet the most demanding service-level requirements.

Microsoft®



Businesses rely on data

Organizations face many challenges today, with greater demands being placed on fewer resources and more pressure than ever to improve performance. Enterprises need to remain competitive by creating new products and services to attract and retain customers, and grow by expanding into new markets and geographies. To stay agile while maintaining growth depends on having ready access to critical business data from across the enterprise. However, current data warehousing systems often struggle under the weight of ever-increasing volumes of data from a wider variety of rapidly-changing sources. The resulting data sprawl, coupled with the increased difficulty of managing complex legacy data warehouse systems, is pushing IT budgets and resources to the breaking point.

Immediate access to relevant, timely business data may often be the most critical factor in this challenging climate, yet many organizations are unable to implement effective large-scale data warehousing solutions for a number of reasons:

- In large enterprises, data is often scattered in disparate places, with no single, consistent view of the business. In addition, a wide variety of data formats, including both structured and unstructured information, make it more difficult to aggregate and analyze data efficiently.
- As enterprises grow, data often ends up locked into increasingly isolated IT systems. Sprawl makes it harder to keep track of all data, even with existing data warehouse solutions, so enterprises end up with an incomplete view of the business.
- Extracting insight from disparate systems often requires manual processes to load data and generate adhoc reports, which increase the burden on already-strained IT resources and decreases IT responsiveness. And ongoing management and optimization of these silos of data is a complicated task, driving up both capital and operational expenditures.

Enterprises looking to succeed in this climate need scalable, highly reliable data warehousing solutions that can overcome these challenges without becoming a burden to IT. Businesses need

- A data warehouse or data mart solution that scales to handle an increasing volume of data from disparate locations, delivering enhanced security and reliability while also ensuring data, governance, and quality to meet regulatory compliance requirements.
- A flexible solution that enables data from mixed workloads across the enterprise to be easily managed without also increasing cost and complexity.

- Automated processes and real-time analytics that draw from all business data to enable faster, better decisions, and deliver a more complete view of the business.
- A fully integrated and tested hardware and software solution that is more cost-effective, less time-consuming, and carries less deployment risk than custom-built systems that require expensive IT resources.

The HP Enterprise Data Warehouse Appliance

HP and Microsoft have come together to help enterprises more easily implement data warehousing technology by designing and delivering a complete solution built on a proven reference architecture and with an optimized configuration. The HP Enterprise Data Warehouse Appliance incorporates the Microsoft application platform and HP Converged Infrastructure to offer an integrated, end-to-end solution that delivers enterprise-class scalability and a single view of the truth across your organization's data.

Built on a tightly integrated architecture, the HP Enterprise Data Warehouse Appliance combines high density, capacity, and performance to provide a very large data warehouse with market-leading cost-per-storage, supporting from tens to hundreds of terabytes (TB) of data. Consisting of one control rack and up-to four data racks, this dedicated appliance is a fully integrated stack incorporating HP ProLiant servers, HP StorageWorks storage, and HP ProCurve networking products; combined with backup tools and management software from both HP and Microsoft, pre-installed and pre-configured in the HP factory to simplify implementation and time-to-business value.

Powered by SQL Server 2008 R2 Parallel Data Warehouse, the HP Enterprise Data Warehouse Appliance offers reliable and massive parallel scale-out query processing without the complexity of deploying extremely large data marts. The appliance is built on "golden images" that have been jointly tested, tuned, and validated by HP and Microsoft labs, with components pre-installed, pre-configured, and validated at the factory. This integrated architectural approach ensures a simplified and accelerated deployment with reduced configuration time and streamlined data loading and data center interoperability.

The HP Enterprise Data Warehouse Appliance offers a balanced system that ensures high-performance and scalability at reduced cost, and with a faster time-to-deployment. For example, an integrated appliance on industry-standard hardware and software can reduce setup, configuration, and expansion time by one to six months on average, thus minimizing implementation costs and risks.

The competitive difference

With the HP Enterprise Data Warehouse appliance, HP and Microsoft are combining their expertise to deliver a solution that offers greater scalability and higher performance than any other SQL Server data warehouse deployment existing today. In addition, the HP Enterprise Data Warehouse appliance brings real and measurable benefits over competitive solutions:

- 25% greater processor scalability, better performance and scalability per data rack¹
- 25% greater scalability of storage capacity per data rack, which lowers costs by supporting larger data warehouses with fewer racks²
- Up to 50% better tempdb performance for complex queries and staging-to-production operations, to achieve performance targets with fewer nodes and better performance and scalability per rack³
- Up to 1.5x superior scan-rate performance by leveraging small form-factor disks for the data nodes, at similar cost-per-terabyte to large form-factor configurations
- Only HP offers the end-to-end value chain from sale and build to support for end-to-end simplicity and support peace-of-mind

^{1,2}Source: HP Engineering data, August 2010

³Source: Based on Oracle public pricing lists: Exadata price list and Oracle database price list. Assumes the following costs: Hardware \$650K (168TB of raw space/~60TB actual data space); Storage software \$1.68M (\$10,000/drive); Database \$2.99M (8 dual processor, quad core servers with 11g, EE+RAC+Partition+Compression)

“Industry-leading organizations effectively leverage business insight to build competitive advantage—IDC research shows that 80% of leaders demonstrate the highest level of analytical orientation enabled by data warehousing and business intelligence solutions.”

—IDC, Benefits of Investing in Data Warehousing to Enable BI and Analytics, July 2010

Solution benefits and features





For enterprises requiring an integrated data warehouse solution with the right price-performance ratio, the HP Enterprise Data Warehouse Appliance offers a powerful, reliable solution that scales cost-effectively to meet their data needs. This appliance makes scalable data warehousing available to more organizations by consolidating separate, dedicated database resources on a standards-based infrastructure that is easy to deploy, operate, and expand.

The HP Enterprise Data Warehouse includes the following benefits and features:

- **Powerful Performance**—Realize best-in-class performance, throughput, and capacity with an optimized architecture that includes ten active compute nodes per rack, with up to 24 SAS disks per node and 8 internal SAS disks (including 6 SAS disks dedicated to tempdb) per server node.
- **Greater Scalability**—Maximize availability and handling of data through a Massively Parallel Processing (MPP) architecture built on carefully selected and optimized components. Gain enterprise-class scalability through distributed query optimization and a unique “Shared Nothing Architecture” that is extensible through a hub-and-spoke model to handle increasing volumes of data.
- **Cost-Effective Operation**—Reduce deployment and operation costs with an integrated appliance architecture that is faster and simpler to implement. Lower acquisition costs and gain greater business value through included Microsoft business intelligence tools (Analysis Services, Integration Services, Reporting Services, SharePoint® Server, and PowerPivot™.) that maximize return on investment by using technology your business users are already familiar with.

- **End-to-End Integration**—Take advantage of a comprehensive platform including Microsoft tools for Extract, Transform, and Load (ETL); Business Intelligence (BI); and Master Data Management (MDM), with streaming data from StreamInsight for a complete data warehouse solution spanning desktop, data marts, and enterprise data warehouses.
- **Appliance-Like Simplicity**—Ease deployment and reduce configuration time through hardware pre-imaged with Windows Server® 2008 R2 and SQL Server 2008 R2 Parallel Data Warehouse software delivering an appliance-like solution that is optimized to install in less than a day.
- **Intelligent Control**—Reduce complexity and streamline ongoing administration through integrated management incorporating HP iLO advanced. Leverage built-in HP innovations for integrated light-out, remote management to maintain health and performance from anywhere to reduce downtime.
- **Reduced Risk**—Reduce time, cost, and uncertainty by incorporating best practices from the HP and Microsoft Fast Track Data Warehouse reference architecture, which dramatically simplifies integration and deployment compared to manually selecting and configuring hardware for a custom solution. Pre-built, optimized golden images and pre-configured hardware delivered by HP, coupled with proven data warehouse best practices, provide consistent, error-free installation in days rather than months.
- **Support Peace-of-Mind**—Gain peace-of-mind with an appliance solution jointly tested and validated by Microsoft and HP. Rely on HP services expertise for data loading, data management, and data center integration. HP and Microsoft jointly deliver industry-leading reactive and mission critical support to provide integrated and end-to-end coverage of both the hardware and software components for your appliance.
- **Future-Proofed Investment**—HP data warehouse solutions are built with growth in mind. Easily leverage your initial investment in a Fast Track Data Warehouse by using it in conjunction with your new HP Enterprise Data Warehouse. You can also grow capacity by adding a new data rack or Fast Track Data Warehouse to your HP Enterprise Data Warehouse appliance.

Solution components

Components		
 	Servers	2x HP ProLiant DL360 G7; 1x DL370 G6; 3x DL380 G7
	Storage	HP StorageWorks MSA P2000 G3 (+ internal) storage
	Software	Microsoft Windows Server 2008 SP2 Microsoft SQL Server 2008 R2 Parallel Data Warehouse*
	Infrastructure	InfiniBand, FC, Ethernet switching, 42U rack
 	Servers	10 active + 1 passive HP ProLiant DL360 G7 compute nodes
	Storage	10x HP StorageWorks MSA P2000 G3
	Software	Microsoft Windows Server 2008 SP2 Microsoft SQL Server 2008 R2 Parallel Data Warehouse*
	Infrastructure	InfiniBand, FC and Ethernet switching, 42U rack
Expansion	Grow from 1–4 data racks, storage options, test/dev system	
Services	Industry-leading reactive and mission critical support jointly delivered by HP and Microsoft to provide integrated and end-to-end coverage of both the hardware and software components of the appliance	

*Software license for Parallel Data Warehouse must be purchased to activate pre-installed software.

HP/Microsoft Converged Applications Appliances

HP and Microsoft are collaborating to optimize and package applications on HP Converged Infrastructure. These applications are delivered through integrated appliances that provide

- **Faster time to value**—Take less time to design, build, and deploy a business intelligence (BI), data warehouse (DW), or online transaction processing (OLTP) solution on the latest Microsoft SQL Server 2008 R2 and HP Converged Infrastructure technology.
- **Optimized performance and scalability**—Achieve higher levels of performance and scalability from pre-tested, optimized solutions that are tuned and balanced to deliver optimal results for specific workloads, and built using the highly scalable SQL Server platform running on high-performance, industry-standard HP server, storage, and networking components.
- **Reduced risk**—HP and Microsoft database appliances are built on proven platforms with built-in high-availability and security, and supported by deep best practices from thousands of successful SQL Server deployments and the commitment of two industry leaders.
- **Enhanced operations and control**—Pre-integrated solutions are easier to buy, install, and maintain; lowering TCO and enabling IT resources to focus on other business priorities.

HP Enterprise Data Warehouse Architecture

The HP Enterprise Data Warehouse Appliance leverages extensive best practices and proven reference architectures to deliver an easy-to-deploy, high-density, high-performance infrastructure for very large data warehouses that scales to over 500TB at a low cost-per-terabyte. Consisting of one control rack and up-to-four data racks, this dedicated, fully integrated appliance combines industry-leading HP server, storage, and networking components with Microsoft application platform software and integrated management.

“The number of data warehouses over 3TB will increase from 36% to 59% by 2012 and will double to 34% for data warehouses greater than 10TB.”

—The Data Warehousing Institute, 2009

Availability and additional information

PLEASE NOTE: Your use of:

The HP Enterprise Data Warehouse Appliance will be available in Q4 calendar year 2010 from your HP/Microsoft select partner. Microsoft SQL Server 2008 R2 Parallel Data Warehouse Edition and Microsoft HPC Pack 2008 with SP1 and SP2 software provided on this server hardware is subject to the terms and conditions of the Volume License program agreement under which you have acquired licenses to these Microsoft products. You may not use the above software without first acquiring valid server and/or client access licenses from Microsoft or its authorized distributor or reseller.

For more information see:

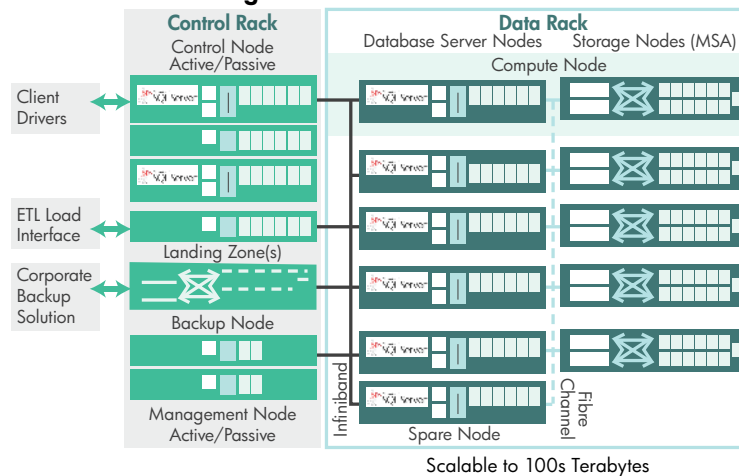
- <http://www.microsoft.com/sqlserver/en/us/get-sql-server/how-to-buy.aspx>

Microsoft Windows Server 2008 Enterprise with SP2 for Data Warehouse Appliances software provided on this server hardware is subject to the Microsoft Software License Terms provided to you by Hewlett-Packard Company or its authorized reseller. The genuine Certificate of Authenticity label affixed to the server hardware identifies licensed software.

See the following website for more information, and learn how you can put the Parallel Data Warehouse Appliance to work for your organization today.

- <http://www.hp.com/solutions/microsoft/edw>

High-Level PDW Architecture



Share with colleagues



Get connected

www.hp.com/go/getconnected

Get the insider view on tech trends, alerts, and HP solutions for better business outcomes

© Copyright 2010 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Microsoft is a U.S. registered trademark of Microsoft Corporation.

4AA2-0987ENW, Created November 2010

Microsoft

