The Hitachi Object Relational (THOR) database technology delivers the power and scaleability you need to overcome the data warehousing challenge. Optimized for the complex query environment, its standard SQL APIs give you the power of a parallel object architecture transparently, using familiar tools. And whether your needs are for decision support, data mining, online analytical processing…text, images, video…gigabytes to terabytes, THOR™ scales effortlessly from two processors to hundreds, growing as your application grows.

Building Your Data Warehouse

You've surmounted the challenge of OLTP. Now you need to interpret the flood of data your mission-critical applications are delivering—to make time-to-market decisions. You need scaleability: your data warehouse may eventually grow to terabytes in size, but already your present RDBMS is gasping for breath, trying to support the demands of your decision support and executive information systems. You need to be able to use familiar client/server tools to deal with new types of data—graphics, maps, video among them—that don't fit the relational mold. And you need to do it all on a budget that seems spread thinner every year.

THOR™ was designed to fit your needs. The Hitachi Object Relational database system is a highly scalable object/relational database accessible through standard ODBC and SQL APIs. You can access THOR's power using familiar client/server tools such as BrioQuery, PowerBuilder, and Visual BASIC, giving you an easy, low-risk migration path. Exploiting massively parallel shared-nothing hardware, THOR's unique data-driven processing effortlessly parallelizes even the most complex queries, and is inherently object-oriented, enabling it to support complex datatypes. The THOR MPP Data Server™ scales almost linearly from four to hundreds of processors, from gigabytes to terabytes of data, yet this power is based on industry-standard, off-the-shelf hardware technology, giving THOR superior price/performance and the ability to exploit industry advances without fear of obsolescence.

Reliability, Accessibility, and Service-ability are built into THOR by Hitachi Computer Products of America (HICAM), backed by the R&D, technological leadership, and financial stability of its $76 billion parent company, Hitachi Ltd., which holds over 85,000 patents or intellectual property rights worldwide. Hitachi is a major provider of large-enterprise solutions to Global Fortune 500 companies. Now, with THOR, Hitachi Computer Products provides the next-generation solution for DSS, OLAP, and other complex query applications.

THOR — Unleashes the Power Within Your Data

The mythical god Thor was armed with magical tools to provide him power and strength. THOR technology provides the ideal solution for the complex queries on very large databases.

- **Superior price/performance** – THOR is based on industry-standard components: the IBM PowerPC chip set, the PCI bus, SCSI, and standard, high-MTBF disk drives.
  - Inherently open, lower-cost solution
  - Industry-proven reliability
  - Exceptional price/performance as your system grows
- **Scalable** – start with as few as four processing elements (each with its own CPU, RAM, disk, and I/O bus) to over 288.
  - Low entry-point to high-end MPP
  - Grows with your applications
  - Near-linear scaleability delivers a consistent and highly favorable ROI on added processing modules
- **Easy, low-risk migration** – you access the power of THOR with familiar tools through standard ODBC and SQL APIs. THOR supports Sybase Open Client/Server.
  - Greatly reduces training costs
  - Quicker proof of concept
  - Many third-party migration aids
  - Compatible with all tools that use ODBC and Sybase Open Server
- **Efficient parallel operation** – THOR's unique data-driven processing, where independent nodes work on operands as they become available, provides an inherent parallelism not matched by other MPP implementations
  - Toroidal mesh interconnect mini-mizes hops between nodes for faster communications, near linear scaleability
  - High speed communication processors offload the CPUs for lower overhead and maximum efficiency
  - Serialization eliminates the need for row or page locking while maintaining ACID properties
  - State of the art techniques such as parallel data access and parallel aggregation ensure high performance
- **Inherently object-oriented** – THOR handles all data, even traditional relational data, in the form of objects that encapsulate data and its related rules, triggers, and other operations, making it easy to support complex datatypes.
  - THOR SQL objectManager™ exploits inherent parallelism of objects for maximum query efficiency
  - Automatically distributes queries, rows, and other objects among processing elements for efficient retrieval
  - Applications can reference objects easily and reliably, including new media types like audio and video.
- **Reliable, Accessible, and Serviceable** – industry standard components and interconnects mean simplicity, reliability, and ease-of-service
  - Mirrored data and transparent backup/recovery prevent reloading data in case of failure
  - Parallel backup and restore activities minimize down time

The competitive power of DSS, EIS, OLAP, and other complex query applications is no longer in doubt—the average payback time for such systems has been estimated at 18-24 months. But the data warehouses that underlie these systems have very different processing requirements from the OLTP-optimized databases that have been a major focus of RDBMS technology for the past decade.

There are three fundamental differences:

- Data warehouses are often far larger than the data sets common in OLTP, and the data to be analyzed are growing faster than processor power is increasing. Shared memory parallel (SMP) systems that easily handle large OLTP applications cannot scale fast enough to grow as your data warehouse applications will.
- Complex queries are much harder to parallelize than the simple atomic transactions of OLTP; parallel systems engineered for OLTP can falter under the load imposed by complex queries.
- Data warehouses must increasingly house data that doesn't fit the relational mode: not only various kinds of multimedia, but also complex logical constructs like invoices or parts explosions.

THOR delivers performance optimized for complex query processing using three synergistic technologies: massively parallel processing (MPP) on shared-nothing hardware, data-driven processing, and object/oriented constructs. Yet all of these powerful techniques are accessible through standard client/server tools: THOR presents a single logical database to clients.

The inherent scaleability of THOR's MPP hardware enables it to start small, yet scale well beyond any SMP database. MPP shared-nothing systems scale linearly; they are not limited by the interconnect performance, I/O, and locking bottlenecks of SMP and shared-disk systems. Instead, each processing node is a completely self-contained computer, with its own CPU, RAM, SCSI bus, and disk storage, communicating with other nodes via a mesh of point-to-point connections.

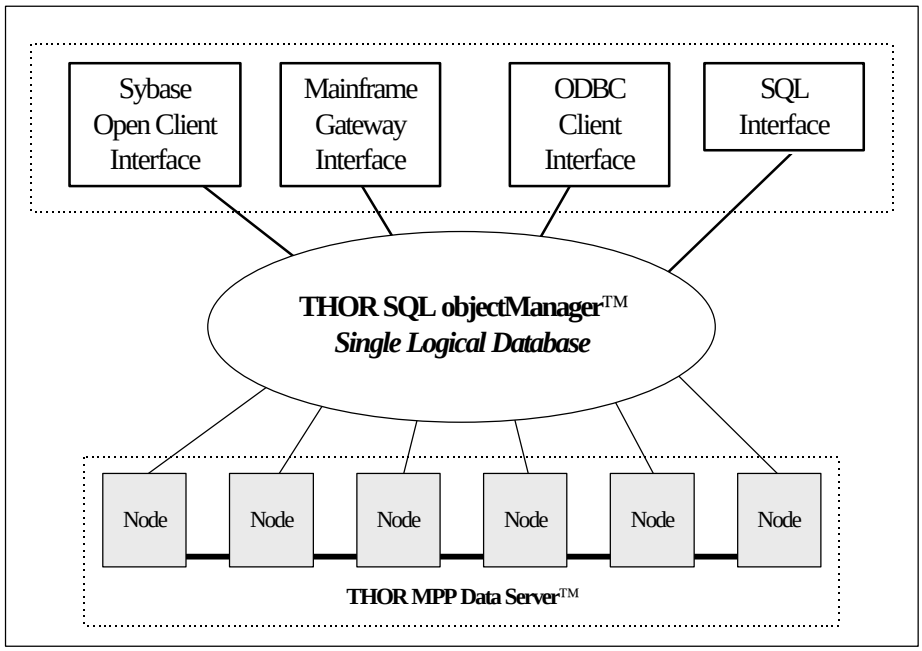

THOR's data-driven processing makes efficient use of its parallel hardware by enabling these nodes to asynchronously receive data from each other, operate on it, and pass it to the next node, in effect pipelining the query in a wave of data passing through THOR's highly optimized toroidal mesh architecture. THOR's SQL objectManager automatically decomposes and assigns the components of complex queries to each node, which operate independently on the data, yielding the maximum possible parallelization without execution plan tuning or manual partitioning.

Of all MPP implementations, only THOR's SQL objectManager is inherently object-oriented: all data is encapsulated with the behaviors or functions needed to operate on it. Not only is object-orientation an efficient software fit to the independence of THOR's shared-nothing architecture, but it makes THOR naturally extensible to support new datatypes.

THOR Tower configuration

The THOR system consists of the THOR SQL object Manager software running on the THOR MPP Data Server hardware. Because THOR runs as a database server to a standard Sybase Open Server supporting standard clients, users need see no change in their applications.

Based on open standards, SQL objectManager is designed to be easily portable to other MPP shared-nothing platforms and will be featured on the most popular of these to meet customer demands.

THOR SQL objectManager™

- Complete SQL implementation optimized for massively parallel shared-nothing platforms
- Data-driven processing delivers superior parallelization of object operations and queries
- Guaranteed processing integrity with transaction management and concurrency control
- Standard API interfaces: Sybase Open Server and Sybase Transact-SQL, and ODBC
- Fundamentally object-oriented design efficiently exploits MPP hardware for both traditional relational data and new datatypes

THOR MPP Data Server™

- Massively parallel shared-nothing architecture delivers most efficient data processing possible
- Based on powerful, off-the-shelf components for cost-effectiveness and reliability
- Modular design for easy scaleability
- Toroidal mesh architecture delivers more cost-effective parallelism than multi-stage or crossbar switches
- Toroidal mesh provides built-in redundancy and high system availability
- Operating system works with SQL objectManager to automatically distribute user data, manage each node's devices, process user workloads, and provide communications, enabling processing power and database capacity to scale with near linearity

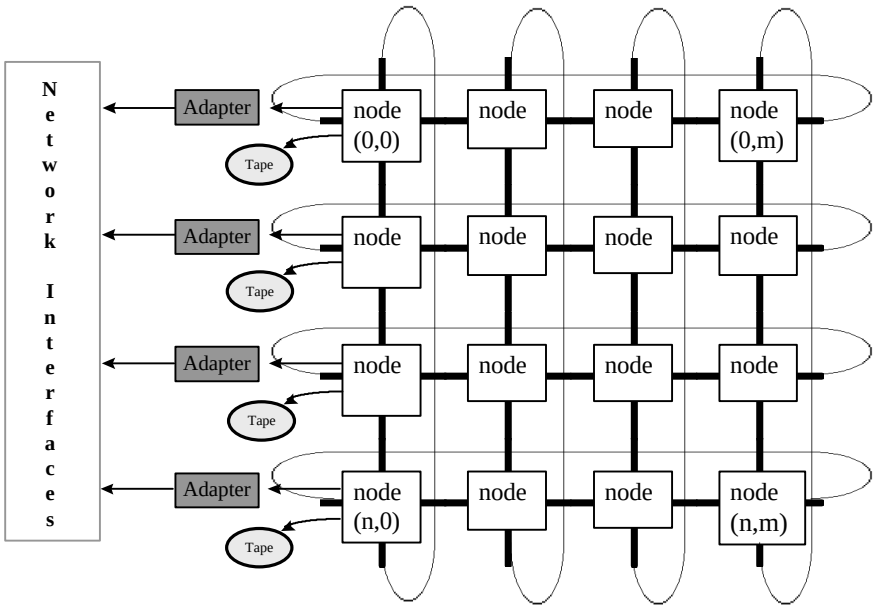

Toroidal Mesh Architecture

The geometry of THOR's toroidal mesh architecture promotes better resource sharing than higher-dimensional networks, providing more scaleability, efficiency, cost effectiveness, and redundancy.

Specifications

Parallel Everything

Product/ Feature	THOR™ Implementation
Data partitioning	Hash
Indexing	Hash
Isolation	Yes
Parallel-Scan	Yes
Parallel-Sort	Yes
Parallel-Aggregate	Yes
Parallel-Join	Yes
Parallel-Import	Yes
Parallel-Dump	Yes
Parallel-Load	Yes
Parallel-Recovery	Yes
Parallel-Insert	Yes
Parallel-Delete	Yes
Disk mirroring	Yes

THOR Technology is World-Class

- Shared-nothing architecture
- SQL-compliant command set
- Open interface: Sybase Open Server
- Server-enforced integrity and security
  - Row level concurrency control
  - Stored procedures and triggers
  - Transaction logging and recovery
  - Database backup and recovery
  - Database administration utilities
- Inherently object-oriented parallel
- Off the shelf – for reliability and cost efficiency
- All operations are encapsulated objects
- New data type support

Hitachi is working closely with beta partners to refine the THOR Database System technology. If you want more information on THOR hardware and software, please contact us in any of the following ways:

Telephone: 1-800-588-THOR
(1-800-588-8467)
408-588-3300

FAX: 408-988-1279

The World Wide Web:
http://www.hicam.hitachi.com/thor

Hitachi Computer Products (America), Inc., a subsidiary of Hitachi America, Ltd., (HICAM) manufactures a wide range of mainframe computer peripherals including disk and magnetic tape drives and semiconductor storage devices. HICAM also develops software and hardware products for networking.