

# MANAGING A DATA WAREHOUSE

Richard Barker  
Chief Technology Officer, SVP  
Veritas Software Corporation, Chertsey, United Kingdom

## Abstract

Major corporations around the world have come to rely totally upon the essential information asset stored in their corporate databases and data warehouses. If the information is not available for *any* reason the business may grind to a halt, and if it stays unavailable for a protracted period, serious financial consequences may result. Making a data warehouse available is not easy. Corporate data warehouses range from one terabyte to ten terabytes or more, with the intention of giving users global access on a 24x365 basis. Data warehouses can take months to set up, yet can fail in seconds. And to react to changing business requirements the data warehouse will need to change in design, content and physical characteristics on a timely basis.

This paper asks a few simple questions:

- ◆ How can I ensure that my data warehouse is constantly available to its users?
- ◆ What happens if the operating system, or disk, or database crashes? – Or the site gets burned down?
- ◆ How can I tune the system on the fly?
  - In fact, how can I change the hardware configuration on the fly?
- ◆ How can I spread the performance load across my global wide-area network?
- ◆ How do I manage such an enormous amount of online and offline critical data?

The paper describes the characteristics of data warehouses and how to make a data warehouse safe, available, perform well, and be manageable. The presentation will cover various means of taking hot backups of huge data warehouses and their business-timely restoration. The exploitation of local and wide-area network replication will be shown to improve performance and availability of global services. New clustering technology will be shown to enable a data warehouse to grow organically to meet the need, exploiting ever changing configurations of storage, clustered servers, Storage Area Networks (SAN) and Local/Wide Area Networks. Other topics considered are disaster recovery, normal high availability, and how we might further *assure the protection and availability of your corporate asset – your corporate data warehouse.*

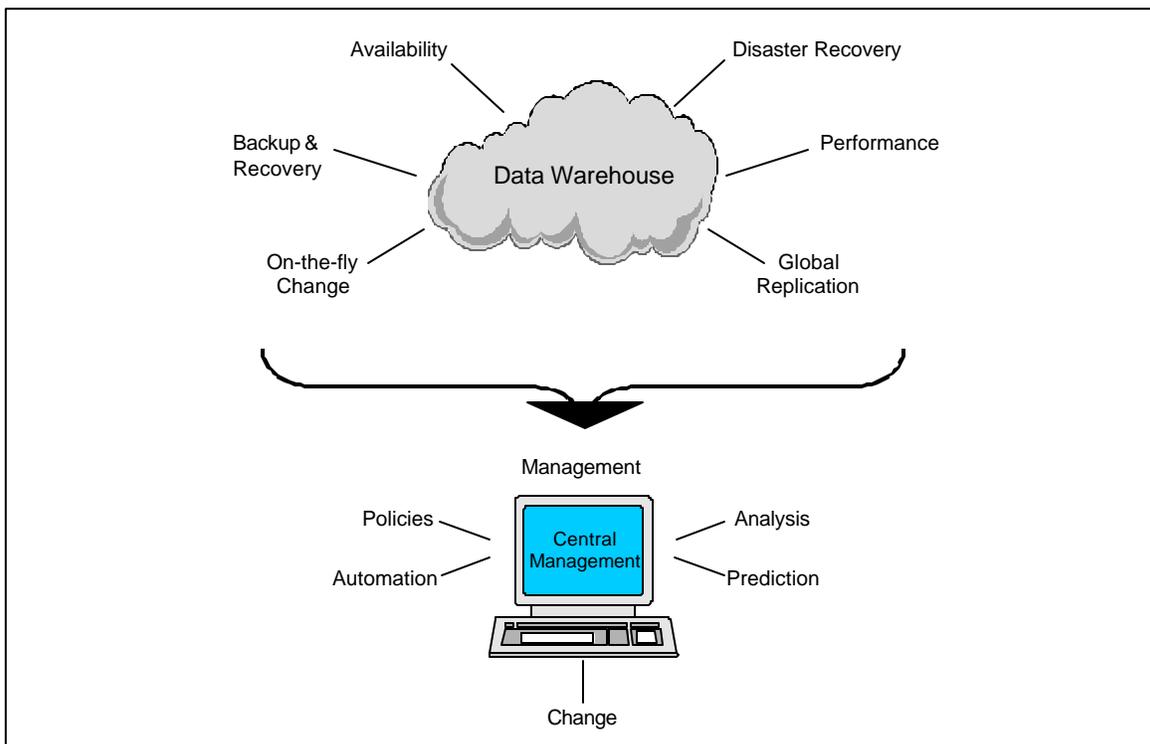
## MANAGING A DATA WAREHOUSE

### What is a Data Warehouse?

Companies set up data warehouses when it is perceived that a body of data is critical to the successful running of their business. Such data may come from a wide variety of sources, and is then typically made available via a coherent database mechanism, such as an Oracle database. By their very nature they tend to be very large in size, used by a large percentage of key employees, and may need to be accessible across the nation, or even the globe. The perceived value of a data warehouse is that executives and others can gain real competitive advantage by having instant access to relevant corporate information. In fact, leading companies now realize that information has become the *most potent* corporate asset in today's demanding markets. Depending on the business, a data warehouse may contain very different things, ranging from the traditional financial, manufacturing, order and customer data, through document, legal and project data, on to the brave new world of market data, press, multi-media, and links to Internet and Intranet web sites. They all have a common set of characteristics: sheer size, interrelated data from many sources, and access by hundreds and thousands of employees.

In placing your organization's essential data in such a single central location, care must be taken to ensure that the information contained within it is highly available, easily managed, secure, backed up and recoverable – to any point in time – and that its performance meets the demanding needs of the modern user.

Figure 1: Managing a Global 24 x 365 Data Warehouse



Within a short time of deployment, a well designed data warehouse will become part of the life blood of its organization. More than that, the company will soon become dependent upon the availability of its data warehouse – placing extreme pressure on the corporate computing department to keep the system going ... and going ... and going.

## Data Warehouse Components

In most cases the data warehouse will have been created by merging related data from many different sources into a single database – a copy-managed data warehouse as in Figure 2. More sophisticated systems also copy related files that may be better kept outside the database for such things as graphs, drawings, word processing documents, images, sound, and so on. Further, other files or data sources may be accessed by links back to the original source, across to special Intranet sites or out to Internet or partner sites. There is often a mixture of current, instantaneous and unfiltered data alongside more structured data. The latter is often summarized and coherent information as it might relate to a quarterly period or a snapshot of the business as of close of day. In each case the goal is better information made available quickly to the decision makers to enable them to get to market faster, drive revenue and service levels, and manage business change.

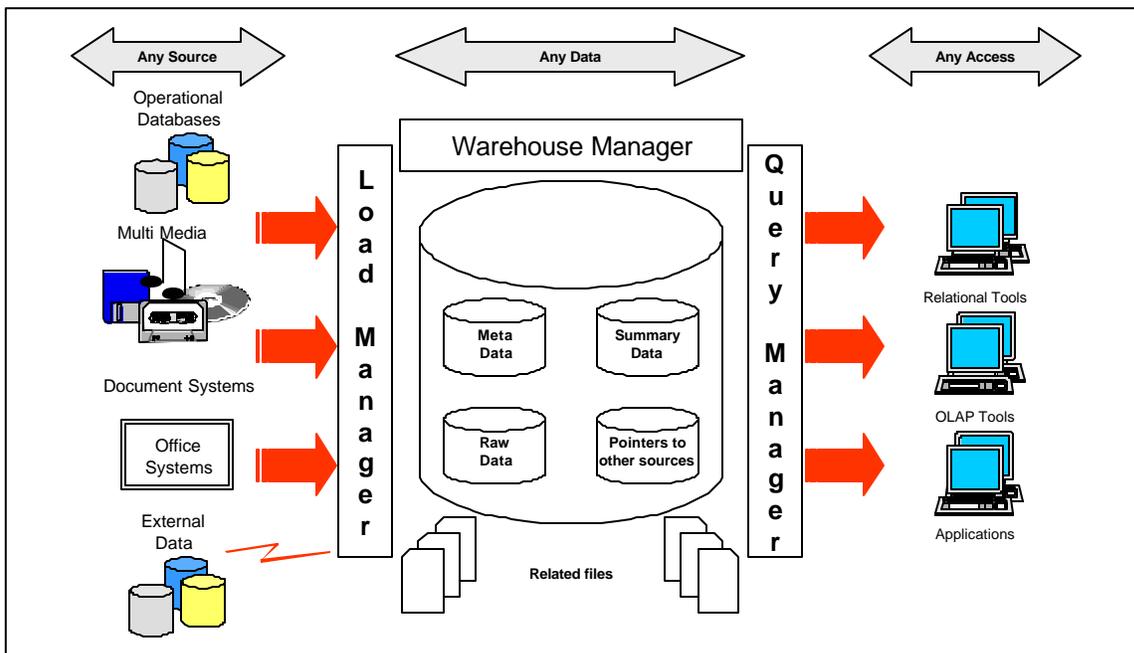
A data warehouse typically has three parts – a load management, a warehouse, and a query management component.

**LOAD MANAGEMENT** relates to the collection of information from disparate internal or external sources. In most cases the loading process includes summarizing, manipulating and changing the data structures into a format that lends itself to analytical processing. Actual raw data should be kept alongside, or within, the data warehouse itself, thereby enabling the construction of new and different representations. A worst-case scenario, if the raw data is not stored, would be to reassemble the data from the various disparate sources around the organization simply to facilitate a different analysis.

**WAREHOUSE MANAGEMENT** relates to the day-to-day management of the data warehouse. The management tasks associated with the warehouse include ensuring its availability, the effective backup of its contents, and its security.

**QUERY MANAGEMENT** relates to the provision of access to the contents of the warehouse and may include the partitioning of information into different areas with different privileges to different users. Access may be provided through custom-built applications, or ad hoc query tools.

Figure 2: A Typical Data Warehouse



## Setting up a Data Warehouse

Such a major undertaking must be considered very carefully. Most companies go through a process like Figure 3 below.

*BUSINESS ANALYSIS*: the determination of what information is required to manage the business competitively.

*DATA ANALYSIS*: the examination of existing systems to determine what is available and identify inconsistencies for resolution.

*SYNTHESIS OF MODEL*: the evolution of a new corporate information model, including meta data – data about data.

*DESIGN AND PREDICT*: the design of the physical data warehouse, its creation and update mechanism, the means of presentation to the user and any supporting infrastructure. In addition, estimation of the size of the data warehouse, growth factors, throughput and response times, and the elapsed time and resources required to create or update the data warehouse.

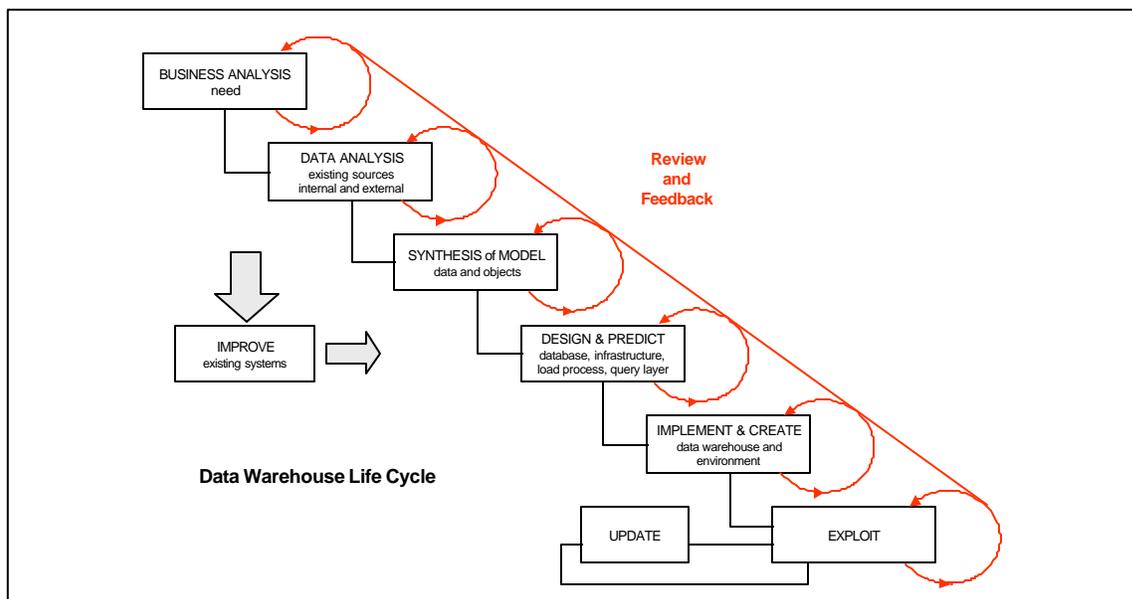
*IMPLEMENT AND CREATE*: the building or generation of new software components to meet the need, along with prototypes, trials, high-volume stress testing, the initial creation of the data, and resolution of inconsistencies.

*EXPLOIT*: the usage and exploitation of the data warehouse.

*UPDATE*: the updating of the data warehouse with new data. This may involve a mix of monthly, weekly, daily, hourly and instantaneous updates of data and links to various data sources.

A key aspect of such a process is a feedback loop to improve or replace existing data sources and to refine the data warehouse given the changing market and business. A part that is often given little focus is to ensure that the system stays operating to the required service levels to provide the maximum business continuity.

Figure 3: Data Warehouse Life Cycle



## Keeping Your Data Warehouse Safe and Available

It can take six months or more to create a data warehouse, but only a few minutes to lose it! Accidents happen. Planning is essential. A well planned operation has fewer ‘accidents’, and when they occur recovery is far more controlled and timely.

### **Backup and Restore**

The fundamental level of safety that must be put in place is a backup system that automates the process and guarantees that the database can be restored with full data integrity in a timely manner. The first step is to ensure that *all* of the *data sources* from which the data warehouse is created are themselves backed up. Even a small file that is used to help integrate larger data sources may play a critical part. Where a data source is external it may be expedient to ‘cache’ the data to disk, to be able to back it up as well. Then there is the requirement to produce say a weekly backup of the entire warehouse itself which can be restored as a coherent whole with full data integrity. Amazingly many companies do not attempt this. They rely on a mirrored system not failing, or recreating the warehouse from scratch. Guess what? They do not even practise the recreation process, so *when* (not if) the system breaks the business impact will be enormous. Backing up the data warehouse itself is *fundamental*.

What must we back up? First the database itself, but also any other files or links that are a key part of its operation. How do we back it up? The simplest answer is to quiesce the entire data warehouse and do a ‘cold’ backup of the database and related files. This is often not an option as they may need to be operational on a nonstop basis, and even if they can be stopped there may not be a large enough window to do the backup. The preferred solution is to do a ‘hot’ database backup – that is, back up the database *and* the related files while they are being updated. This requires a high-end backup product that is synchronized with the database system’s own recovery system *and* has a ‘hot-file’ backup capability to be able to back up the conventional file system. Veritas supports a range of alternative ways of backing up and recovering a data warehouse – for this paper we will consider this to be a very large Oracle 7 or 8 database with a huge number of related files.

The first mechanism is simply to take cold backups of the whole environment, exploiting multiplexing and other techniques to minimize the backup window (or restore time) by exploiting to the full the speed and capacity of the many types and instances of tape and robotics devices that may need to be configured.

The second method is to use the standard interfaces provided by Oracle (Sybase, Informix, SQL BackTrack, SQL server, etc.) to synchronize a backup of the database with the RDBMS recovery mechanism to provide a simple level of ‘hot’ backup of the database, concurrently with any related files. Note that a ‘hot file system’ or checkpointing facility is also used to assure the conventional files backed up correspond to the database.

The third mechanism is to exploit the RDBMS special hot backup mechanisms provided by Oracle and others. The responsibility for the database part of the data warehouse is taken by, say, Oracle, who provide a set of data streams to the backup system – and later request parts back for restore purposes. With Oracle 7 and 8 this can be used for very fast full backups, and the Veritas NetBackup facility can again ensure that other non-database files are backed up to cover the whole data warehouse. Oracle 8 can also be used with the Veritas NetBackup product to take incremental backups of the database. Each backup, however, requires Oracle to take a scan of the entire data warehouse to determine which blocks have changed, prior to providing the data stream to the backup process. This could be a severe overhead on a 5 terabyte data warehouse. These mechanisms can be fine tuned by partition etc. Also optimizations can be done – for example, to back up read-only partitions once only (or occasionally) and to optionally not back up indexes that can easily be recreated.

Veritas now uniquely also supports block-level incremental backup of any database or file system without requiring pre-scanning. This exploits file-system-level storage checkpoints. The facility will also be available with the notion of ‘synthetic full’ backups where the last full backup can be merged with a set of incremental backups (or the last cumulative incremental backup) to create a new full backup *off line*. These two facilities reduce by several orders of magnitude the time and resources taken to back up and restore a large data

warehouse. Veritas also supports the notion of storage checkpoint recovery by which means the data warehouse can be instantly reset to a date/time when a checkpoint was taken; for example, 7.00 a.m. on each working day.

The technology can be integrated with the Veritas Volume Management technology to automate the taking of full backups of a data warehouse by means of third-mirror breakoff and backup of that mirror. This is particularly relevant for smaller data warehouses where having a second and third copy of it can be exploited for resilience and backup. Replication technology, at either the volume or full-system level, can also be used to keep an up-to-date copy of the data warehouse on a local or remote site – another form of instantaneous, always available, full backup of the data warehouse.

And finally, the Veritas software can be used to exploit network-attached intelligent disk and tape arrays to take backups of the data warehouse directly from disk to tape – without going through the server technology. Alternatives here are disk to tape on a single, intelligent, network-attached device, or disk to tape across a fiber channel from one network-attached device to another.

Each of these backup and restore methods addresses different service-level needs for the data warehouse and also the particular type and number of computers, offline devices and network configurations that may be available. In many corporations a hybrid or combination may be employed to balance cost and services levels.

### **Online Versus Offline Storage**

With the growth in data usage – often more than doubling each year – it is important to assure that the data warehouse can utilize the correct balance of offline as well as online storage. A well balanced system can help control the growth and avoid ‘disk full’ problems, which cause more than 20% of stoppages on big complex systems. Candidates for offline storage include old raw data, old reports, and rarely used multi media and documents.

Hierarchical Storage Management (HSM) is the ability to ‘off line’ files automatically to secondary storage, yet leaving them accessible to the user. The user sees the file and can access it, but is actually looking at a small ‘stub’ since the bulk of the file has been moved elsewhere. When accessed, the file is returned to the online storage and manipulated by the user with only a small delay. The significance of this to a data warehousing environment in the first instance relates to user activities around the data warehouse. Generally speaking, users will access the data warehouse and run reports of varying sophistication. The output from these will either be viewed dynamically on screen or held on a file server. Old reports are useful for comparative purposes, but are infrequently accessed and can consume huge quantities of disk space. The Veritas HSM system provides an effective way to manage these disk-space problems by migrating files, of any particular type, to secondary or tertiary storage. Perhaps the largest benefit is to migrate off line the truly enormous amounts of old raw data sources, leaving them ‘apparently’ on line in case they are needed again for some new critical analysis. Veritas HSM and NetBackup are tightly integrated, which provides another immediate benefit – reduced backup times – since the backup is now simply of ‘stubs’ instead of complete files.

### **Disaster Recovery**

From a business perspective the next most important thing may well be to have a disaster recovery site set up to which copies of all the key systems are sent regularly. Several techniques can be used, ranging from manual copies to full automation. The simplest mechanism is to use a backup product that can automatically produce copies for a remote site. For more complex environments, particularly where there is a hybrid of database management systems and conventional files, an HSM system can be used to copy data to a disaster recovery site automatically. Policies can be set so files or backup files can be migrated automatically from media type to media type, and from site to site; for example, disk to optical, to tape, to an off-site vault.

Where companies can afford the redundant hardware and very-high-bandwidth (fiber channel) communications wide-area network, volume replication can be used to retain a secondary remote site *identical* to the primary data warehouse site.

### **Reliability and High Availability**

A reliable data warehouse needs to depend upon restore and recovery a lot less. After choosing reliable hardware and software the most obvious thing is to use redundant disk technology, dramatically improving both reliability and performance, which are often key in measuring end-user availability. Most data warehouses have kept the database on raw partitions rather than on top of a file system – purely to gain performance. The Veritas file system has been given a special mechanism to run the database at exactly the same speed as the raw partitions. The Veritas file system is a journaling file system and recovers in seconds, should there be a crash. This then enables the database administrator to get all the usability benefits of a file system with no loss of performance. When used with the Veritas Volume Manager we also get the benefit of ‘software RAID’, to provide redundant disk-based reliability and performance.

After disks, the next most sensible way of increasing reliability is to use a redundant computer, along with event management and high availability software, such as FirstWatch. The event management software should be used to monitor all aspects of the data warehouse – operating system, files, database, and applications. Should a failure occur the software should deal with the event automatically or escalate instantly to an administrator. In the event of a major problem that cannot be fixed, the High Availability software should fail the system over to the secondary computer within a few seconds. Advanced HA solutions enable the secondary machine to be used for other purposes in the meantime, utilizing this otherwise expensive asset.

### **‘On the Fly’ Physical Reconfiguration**

It is outside the scope of this paper to consider how a database schema might evolve and thus how the content of the data warehouse can change logically over time. At the physical level the data may need to be mapped to disk or other online media in different ways over time.

The simplest case is a data warehouse that grows in volume at say 30% per year (during initial growth this may be 100 to 1000% in a year). The system must be able to extend the database, on top of a file system, which itself must be extensible, and on top of volumes which can also continuously add new disks – *all on the fly*. The combination of Oracle on top of the Veritas file system and volume manager uniquely solves this problem. The volume manager, for example, allows ‘hot swapping’ and hot addition of disks, and can manage disks of different characteristics from different vendors simultaneously.

The volume manager provides sophisticated software RAID technology, optionally complementing hardware RAID components for the most critical data sets. RAID can then be used to provide very high levels of reliability and/or stripe for performance. When this unique software combination is used, the layout of the data warehouse to disk can be changed on the fly to remove a performance bottleneck or cater for a new high influx of data. You can even watch the data move and see the performance improve from the graphics monitoring console! A new facility can be used to monitor access to the physical disks and recommend the re-layout that would be most beneficial to reliability, performance or both.

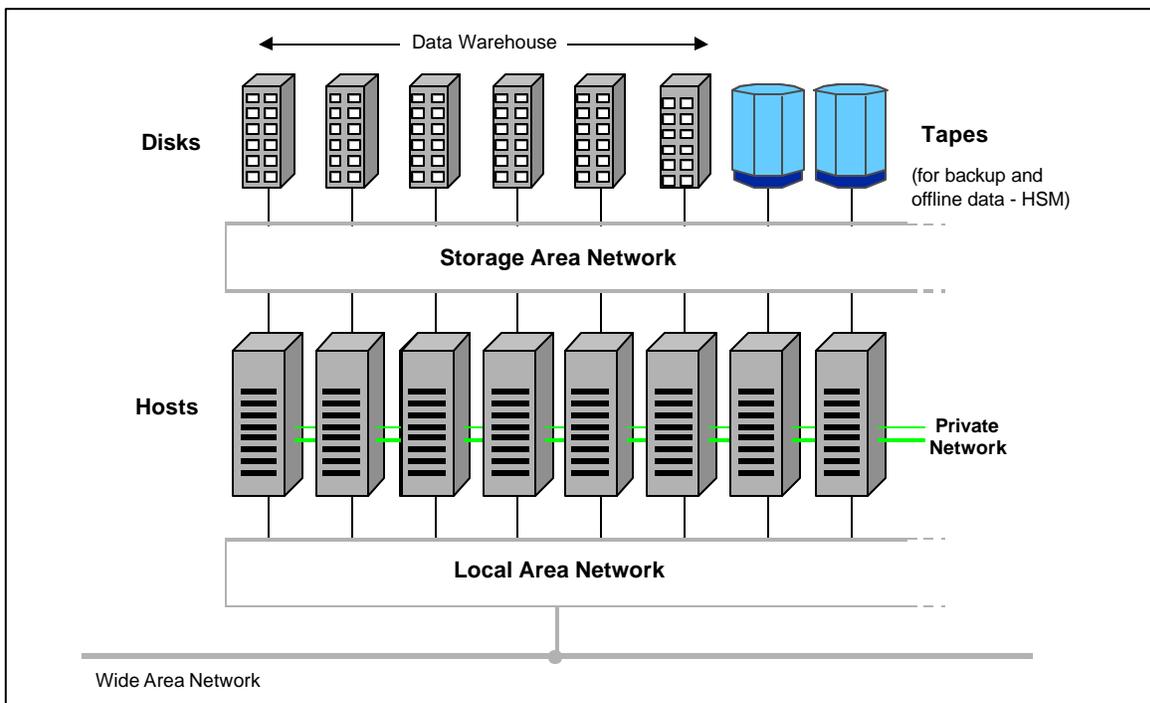
The next most complex case is adding new hosts to provide the horse power to the system. An initial data warehouse may require, say, four powerful computers in a cluster to service the load. Now simple failover technology is less useful and clustered high availability solutions are required to cater for the loss of one or more nodes and the ‘take-over’ of that load by remaining machines. As the system grows the cluster may have to increase from 4 machines to 8, 16, 32 or more. Veritas has developed a new clustering technology which can scale to as many as 500+ computers that collectively supply the horse power to the data warehouse. (Higher numbers may be essential for NT-based solutions.) Within the entire cluster, groups of them can collectively

provide certain functions, like transaction processing, decision support or data mining, and also serve as alternative secondary servers to cater for any host failure in that group. Note: for optimal operation dynamic load balancing may also be required.

This clustering technology enables on-the-fly changes such as applications to be hot-swappable to different servers, and to add or remove servers, add or remove disks or arrays, or otherwise reconfigure the hardware underpinning the data warehouse. *On-the-fly everything*.

To work, the system uses a redundant private network to maintain peer-to-peer control, set and enforce policies, and to instantly accept/drop servers into the cluster. The disk (and tape) arrays may be connected using a fiber channel into a Storage Area Network or connected more conventionally to servers. Access to the data warehouse on the cluster would be via a local area network – possibly again using a fiber channel to provide the best performance and throughput. To complete the picture a redundant site may be maintained, if required, using Wide Area Network replication technologies.

Figure 4: Veritas Advanced Cluster and High Availability Technology



## Replicated Data Warehouses and Disaster Recovery Sites

To manage a global or national data warehouse replication may be very useful for several reasons. From a mission-critical perspective, maintaining a replica on a remote site for disaster recovery can be very important. The ideal mechanism is to use a replicated volume approach, such as the Veritas Replicated Volume Manager. This highly efficient tool can replicate every block written to disk to one or more local or remote sites. Replication may be asynchronous for normal applications (losing a few seconds or minutes of data in the event of a disaster) or synchronous for critical applications where no loss of data whatsoever is acceptable. Note: with replication at the volume level, the replicas cannot be used at the same time as the primary.

Replication can also be useful to off load all read-only applications to a second machine. This is particularly useful for data warehouses, which typically have 80+% of access in a read-only manner – for example, ad hoc

inquiry or data mining. Here the ideal mechanism is the Veritas Replicated File System, which permits the replicas to be read even while they are being synchronously (or asynchronously) maintained from the primary. In a global situation, several replicas could be used – one in each concentrated end-user community – such that all read accesses are done locally for maximum performance to the user. This mechanism provides resilience in the sense of continued local access to the entire warehouse for read access – even when the primary or network is down. Note: this approach uses more server I/O than replicated volumes but offers great functionality.

## Data Warehouse Versioning and Checkpoint Recovery

Veritas has produced an interesting technology whereby one or more checkpoints or clones can be created to give a view of the data warehouse at a particular instant in time – a version. This is achieved by creating a storage checkpoint in the file system that underpins the database and the normal files that constitute the online data warehouse. Subsequent changes of any block in the database continue to be held in the live database, while their corresponding before images of the blocks are held in a ‘checkpoint file’. Thus applications could log on to a storage checkpoint version of the data warehouse, say taken at exactly 6.00 p.m. on a Friday, and conduct a whole series of analysis over several days purely on the warehouse data as it was at that instant. The live database carries on being updated and used normally in parallel.

This facility enables consistent full backups to be taken and supports a new paradigm in decision support and versioning of corporate data warehouses.

The final benefit of this approach is that if the recovery process required is to revert the data warehouse to the moment in time the storage checkpoint was taken, then this can be achieved in a few moments. The ‘storage checkpoint’ acts as a container for block-level before images for the database and can be reapplied very quickly to revert the database to its original state.

## Management Tools

Managing such potential complexity and making decisions about which options to use can become a nightmare. In fact many organizations have had to deploy many more administrators than they had originally planned *just to keep the fires down*. These people cost a lot of money. Being reactive, rather than proactive, means that the resources supporting the data warehouse are not properly deployed – typically resulting in poor performance, excessive numbers of problems, slow reaction to them, and over buying of hardware and software. The ‘when in doubt throw more kit at the problem’ syndrome.

Management tools must therefore address enabling administrators to switch from reactive to proactive management by automating normal and expected exception conditions against policy. The tools must encompass *all* of the storage being managed – which means the database, files, file systems, volumes, disk and tape arrays, intelligent controllers and embedded or other tools that each can manage part of the scene.

To assist the proactive management, the tools must collect data about the data warehouse and how it is (and will be) used. Such data could be high level, such as number of users, size, growth, online/offline mix, access trends from different parts of the world. Or it could be very detailed – such as space capacity on a specific disk or array of disks, access patterns on each of several thousand disks, time to retrieve a multi-media file from offline storage, peak utilization of a server in a cluster, and so on – in other words, raw or aggregate data over time that could be used to help optimize the existing data warehouse configuration. The tools should, ideally, then suggest or recommend better ways of doing things. A simple example would be to analyze disk utilization automatically and recommend moving the backup job to an hour later (or use a different technique) and to stripe the data on a particular partition to improve the performance of the system (without adversely impacting other things – a nicety often forgotten). This style of tool is invaluable when change is required. It can automatically manage and advise on the essential growth of the data warehouse, pre-emptively advise on

problems that will otherwise soon occur using threshold management and trend analysis. The tools can then be used to execute the change envisaged and any subsequent fine tuning that may be required. Once again it is worth noting that the management tools may need to exploit lower-level tools with which it is loosely or tightly integrated. Finally, the data about the data warehouse could be used to produce more accurate capacity models for proactive planning.

Veritas is developing a set of management tools that address these issues. They are:

*STORAGE MANAGER*: managing all storage objects such as the database, file systems, tape and disk arrays, network-attached intelligent devices, etc. The product also automates many administrative processes, manages exceptions, collects data about data, enables online performance monitoring and lets you see the health of the data warehouse at a glance. Storage Manager enables other Veritas and third-party products to be exploited in context, to cover the full spectrum of management required – snap-in tools.

*STORAGE ANALYST*: collects and aggregates further data, and enables analysis of the data over time.

*STORAGE OPTIMIZER*: recommends sensible actions to remove hot spots and otherwise improve the performance or reliability of the online (and later offline) storage based on historical usage patterns.

*STORAGE PLANNER*: will enable capacity planning of online/offline storage, focusing on very large global databases and data warehouses.

[Note: versions of Storage Manager and Optimizer are available now, with the others being phased for later in 1998 and 1999.]

The use of these tools and tools from other vendors should ideally be started during the ‘Design and Predict’ phase of development of a data warehouse. It is, however, an unfortunate truth that in most cases they will have to be used retrospectively to manage situations that have become difficult to control – to regain the initiative with these key corporate assets.

## Conclusions

Data warehouses, datamarts and other large database systems are now critical to most global organizations. Management of them starts with a good life-cycle process that concentrates on the operational aspects of the system. Their success is dependent on the availability, accessibility and performance of the system.

The operational management of a data warehouse should ideally focus on these success factors. Putting the structured part of the database on a leading database, such as Oracle, provides the assurance of the RDBMS vendor and the vendor’s own management tools. Running the database on top of the Veritas File System, along with other data warehouse files, provides maximum ease of management with optimal performance and availability. It also enables the most efficient incremental backup method available when used with the Veritas NetBackup facilities. By exploiting the Veritas Volume Manager the disk arrays can be laid out for the best balance of performance and data resilience. The Optimizer product can identify hot spots and eliminate them on the fly. Replication services at the volume or file-system level can be used to provide enhanced disaster recovery, remote backup and multiple remote sites by which the decision support needs of the data warehouse can be ‘localized’ to the user community – thereby adding further resilience and providing optimal ‘read’ access. High availability and advanced clustering facilities complete the picture of constantly available, growing, high performance data warehouses.

High-level storage management tools can provide a simple view of these sophisticated options, and enable management by policy and exception. They can also add value by analysis of trends, optimization of existing configurations, through to predictive capacity planning of the data warehouse future needs. In summary, the key to the operational success of a global data warehouse is ‘online everything’, where all changes to the software, online and offline storage and the hardware can be done on line on a 24x365 basis. Veritas is *the*

storage company to provide end-to-end storage management, performance and availability for the most ambitious data warehouses of today.

**BIBLIOGRAPHY**

Dr Chris Boorman, *Data Warehouses: Keeping Your Crown Jewels Safe*, EOUG 1998

© Veritas Software Corporation, February 1998, and EOUG 1998