



V
E
R
I
T
A
S

W
H
I
T
E

P
A
P
E
R

Building a Successful E-Business Website

**Using Data Availability Software to
Optimize Site Performance and
Availability**

November 2000

Table of Contents

Introduction	3
What Makes an E-Business Site Successful	4
Steps to Creating a High Performing, Highly Available E-Business Site	5
1. Design the Architecture for Availability	5
2. Choose Highly Available and Manageable Platforms	7
Storage availability	7
Administrative downtime	7
The database layer	7
3. Optimize Performance from Your Systems	8
I/O performance	8
Protecting production systems with off-host processing	9
4. Protect the Data	9
Protection from disk errors	9
Protection from application or user errors	10
Protection from disaster	10
5. Protect the Applications	11
Shared nothing availability clusters	11
Shared data clusters	11
6. Plan for Change	12
Summary	13

Introduction

It is no longer a revelation that the Internet changes everything – just ask the businesses trying to build and maintain successful e-business sites in today's around-the-clock online environment. Maintenance downtime that used to be accepted as a matter of course is now impossible. Usage growth is no longer steady and predictable, but can be sudden and dramatic. User expectations are high, and the consequences of not meeting them are severe.

Although the Internet economy has created many opportunities for new businesses and business models, the online climate has little tolerance for delays or failures. It's not enough any more simply to get a site out there first in your market. E-Business sites must be designed and built for long-term viability from the start.

Whether you are building a new site or making continuous improvements to an existing e-business site, this paper offers guidelines for creating highly available, high performing e-business sites.

What Makes an E-Business Site Successful

Executing your e-business strategy online requires more than careful business planning – it requires the technical expertise to build and maintain a site that can truly support the business over time. You must offer compelling services and these services must be able to handle the unique challenges of the Internet environment.

- **The demand for around-the-clock access.** Consumers expect online services to be available any time of the day or night. If you are in the Business-to-Business market, a global reach requires around-the-clock availability of your systems. There is little tolerance for downtime, whether it is planned or accidental.
- **Performance expectations.** Both business and consumer users have high performance expectations when connecting to your site. Research shows that most users will give up after only a six-second delay, and frequently will not return again. Many analysts attribute the low conversion rates from lookers to buyers on e-business sites to poor site performance – it simply takes too long to get the information they need or complete a transaction.

Zona Research estimates that e-commerce revenue is being lost at the rate of \$362 million per month, and perhaps as much as \$4.35 billion per year, due to failures and unacceptable performance of websites.

- **Unpredictable traffic and growth.** Short-term traffic is hard to predict, surging wildly in response to marketing campaigns, seasonal patterns, or apparently random reasons. Long-term business growth is likewise difficult to plan for in the new economy, defying traditional capacity planning practices. In response to this challenge, an e-business must always plan to handle more than the current traffic load.
- **Rapidly changing market conditions.** Speed to market is driving competition in emerging online markets. With constantly innovating competitors, sites are under pressure to offer increasingly complex processing – more personalization, more interactive services. The e-business site infrastructure must support rapid change and agile business practices.

It is imperative that the e-business builds out an Internet infrastructure that can meet these challenges, providing superior performance and availability and enabling constant growth and change. This paper is concerned with how to successfully *implement* your online strategy in an e-business site, and specifically in the operational aspects of the site.

VERITAS Software, the data availability company, is the ideal partner for your e-business endeavors. VERITAS has strong partnerships with leading system and storage vendors, and years of experience in keeping critical data available. The Fortune 500 companies use VERITAS integrated storage management products to reduce downtime and improve performance for critical systems and databases. Now e-businesses are using integrated solutions from VERITAS to meet the demanding expectations of the online environment. VERITAS, Oracle, and Sun Microsystems have formed a strategic alliance to support e-business growth, providing proven, integrated and tested solutions.

Steps to Creating a High Performing, Highly Available E-Business Site

Whether you are planning to implement a new site or upgrading an existing e-business site to meet growing needs, the following steps provide basic guidelines for meeting the performance and availability demands of today's online business environment.

1. Design an architecture that ensures availability. Build redundancy and scalability in your design so the architecture will grow as your business grows, while being resilient to failures.
2. Choose highly available and manageable system platforms. Make sure that the foundation for the various components is as robust as possible, while creating an operating environment in which you can manage constantly growing storage and accelerating usage.
3. Choose solutions that optimize performance regardless of system load, as the site needs to run at peak performance, all the time.
4. Make sure that the data is protected from hardware component failures and operating system or database corruption.
5. Data is only part of the overall availability question – the application must also be protected, and provisions made for application-specific monitoring and application failover.
6. Plan for change. Nothing is constant in the e-business environment, so choose solutions that are flexible enough to work in dynamic and heterogeneous environments.

The following sections describe these steps in more detail.

1. Design the Architecture for Availability

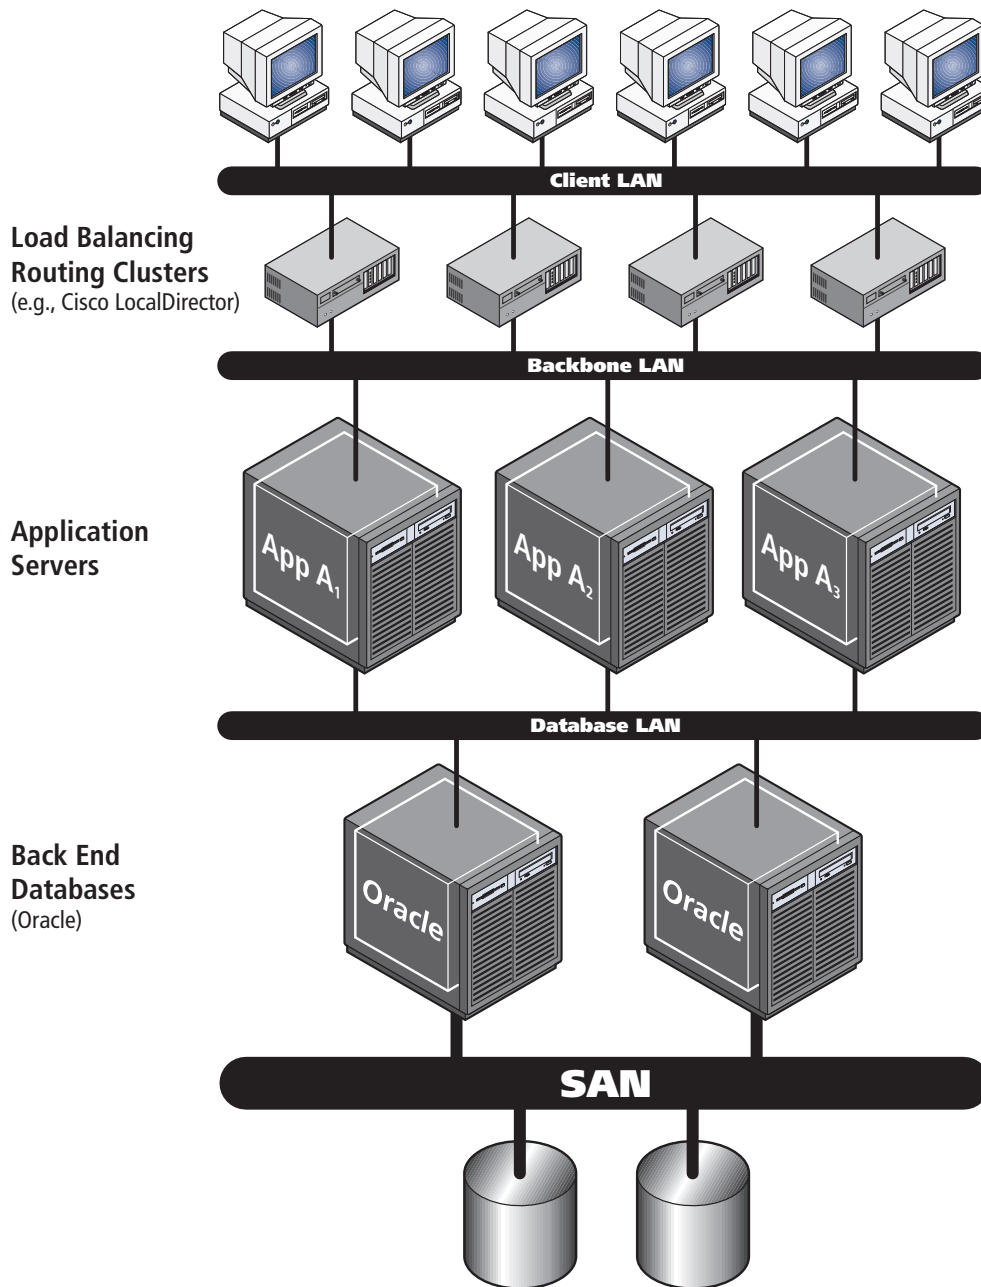
The mature e-business site is multi-tiered, with several hardware components, including:

- Firewalls and routers
- Load balancers
- Web servers/front ends
- Application servers
- Database servers
- File servers
- Storage hardware

The software environment that runs on this hardware can be equally complex, and may include:

- Operating system software
- Web/application server software
- Database software and applications
- Storage management software (backup/recovery, clustering and failover, etc.)

When assembling these components to create your e-business site, it is important to build in redundancy where possible. Redundant equipment can provide both protection from failure and scalable processing for load peaks.



A multi-tiered e-business site has redundancy of equipment at every level.

In the example above, redundant equipment at almost every level provides some degree of protection from failure. For example, storage is shared between database servers; combined with a clustering/failover solution, this enables one server to take over requests for another in case of failure. Likewise, if an application server fails, the load balancers can direct requests to other servers.

Physical hardware redundancy is not enough; clustering software manages an automated failover to clustered equipment when necessary. This is discussed in [Step 5, Protect the Applications](#).

2. Choose Highly Available and Manageable Platforms

All of the critical components of the site should be built on platforms that optimize the availability and manageability of the total solution.

The combined hardware/software platform should keep downtime of all kinds low – including unplanned failures and planned, administrative downtime. Requirements include:

- A high mean time between failures for all hardware components
- Fast recovery from outages when they do occur
- Good administrative tools and policies for system maintenance
- Support for online administrative activities where possible

Storage availability

Critical data must reside on highly available storage. Again, mean time between failures is one important measure. But truly critical data gains protection from disk failures if stored in RAID or mirrored configurations.

RAID (Redundant Arrays of Inexpensive Disk) can be implemented in a storage device or on the host system. Most RAID implementations use striping (for performance) combined with parity checking. A RAID-5 configuration, for example, can withstand the failure of any one device in the array without loss of data. The loss of a second device while before the first is replaced and reloaded does result in data loss, however. For this reason, it is important not to rely solely on the hardware, but also to implement good policies for monitoring and replacing defective hardware.

Administrative downtime

Online administration support is often overlooked, but critically important. In fact, it's a catch-22 for some sites – they need to install the latest software patches, perform frequent backups, defragment storage, and perform other tasks to keep the system performing optimally. But performing these tasks can affect the system availability.

It is important to create an administrative environment that supports administrators – enabling them to optimize the system without causing costly downtime. There are several options for supporting this.

Third mirror split solutions can create a copy of production data for administrative processing purposes. (In this situation, you mirror the production storage, and then split the mirror. If the data is already mirrored, then you create a third mirror for this purpose.) Splitting the mirror creates a copy that can be used for data-intensive activities, such as backups, reporting and testing. Mounting this data from a different host system creates completely "off-host" processing for administrative purposes.

Another option is using clustering software that lets you "fail over" production applications to a different host in a cluster. In this way, you can perform operations like system upgrades while still providing key services.

The database layer

Databases introduce their own stringent administrative activities, typically requiring the specialized expertise of a database administrator (DBA). Database administration tasks including defragmenting storage, I/O tuning, and regular backups. Again, it is important to put processes into place to support these activities, without impacting production system performance.

VERITAS solutions

VERITAS Software provides file system and volume management solutions that turn standard servers into highly available, scalable and manageable platforms for the critical components of your e-business.

[VERITAS Database Edition™ for Oracle®](#)

The Database Edition *for Oracle* offers a highly available and manageable platform for Oracle applications running on e-business web sites. Database Edition *for Oracle* consists of VERITAS Software's industry-leading file system and volume management solutions. Tightly integrated, these products enhance availability of critical data by:

- Creating “virtual” storage volumes that can survive disk or I/O bus failure, using mirroring and RAID techniques.
- Preventing planned administrative downtime with online defragmentation, resizing, tuning, and other storage operations.
- Simplifying administration with centralized management of storage across systems and site components.
- Delivering raw partition performance with the ease of administration of file systems using VERITAS Quick I/O™.
- Allowing fast recovery from logical errors such as database corruptions or dropped tablespaces using Storage Rollback.

[VERITAS Database Replication Option for Oracle®](#)

The VERITAS Database Replication Option *for Oracle* lets you copy and relocate critical production databases for off-host processing purposes. For example, you can use a secondary site to perform testing, backups, decision support extracts, reporting, or training. This protects the performance and availability of the production system by off-loading data-intensive activities.

[VERITAS Database Edition™/HA for Oracle®](#)

The VERITAS Database Edition/HA *for Oracle* adds comprehensive application availability management for both hardware and software resources in clustered server configurations. Using VERITAS Cluster Server™, Database Edition/HA supports clusters of up to 32 servers in both Storage Area Network and traditional environments and protects Oracle applications from system and application failure.

3. Optimize Performance from Your Systems

Many factors affect site performance, from architecture to application design to processor speeds.

Good application design is the most important contributor to performance. After that, processor speed is significant. Building on this groundwork, you can achieve better performance by optimizing I/O and off-loading data-intensive activities from production systems.

I/O performance

For data-intensive applications, I/O performance is a critical component of overall application performance.

Buying high speed storage devices is just one part of the I/O performance puzzle. Other activities you can do to create a high performance I/O platform include:

- Tuning I/O to alleviate hotspots.
- Implementing striping or RAID techniques to distribute I/O among many devices. Striping provides better performance for large, contiguous reads, and for systems with heavy I/O loads.

Protecting production systems with off-host processing

Many necessary administrative tasks, such as backups, reporting, and decision support extracts, can have a significant performance impact on production systems. Creating an off-host processing environment lets you move these tasks to other systems with minimal impact on the production system.

VERITAS products for improving performance

Using [VERITAS Database Edition™ for Oracle®](#), you can achieve better performance from any storage devices, and can change storage configurations dynamically as usage patterns evolve.

- Quick I/O and Cached Quick I/O in VERITAS Database Edition *for Oracle* provide better than raw disk performance using file system storage.
- VERITAS Volume Manager™ provides flexible striping and RAID configurations, distributing I/O load among multiple physical devices for faster performance.
- VERITAS File System™ uses an extent-based allocation scheme that speeds overall throughput.
- VERITAS Volume Manager and VERITAS File System work together to tune file system parameters automatically according to the volume layout.
- The graphical user interface gives administrators an easy way to spot potential I/O hotspots and reallocate storage to optimize performance.

“VERITAS Database Edition for Oracle enabled us to exceed our performance and up time requirements, while providing us the flexibility to scale our solution as our business grows.”

— Eric Anderson, Manager of Systems Operations, Excite@Home

4. Protect the Data

The e-business site both hosts and generates a rich array of data that is a critical asset to the business: web content, user profiles, e-business transaction data, and clickstream data. It is critical to protect this data against loss from a variety of potential risks, including disk failure, application or user error, and physical damage due to localized disasters.

Protection from disk errors

Mirroring and RAID storage techniques can protect data from disk failures, ensuring that data remains available if a single disk device fails. These storage techniques may be implemented in the storage hardware itself (controller-based RAID and mirroring) or in software residing on the host system (host-based RAID and mirroring).

RAID-5 storage combines striping with distributed parity, protecting data against the loss of one disk in the array and improving performance. When a RAID device fails, the RAID software can reconstruct the data on the failed device. It is imperative, however, to replace the failed device as quickly as possible, as another disk failure on the same device would lead to data loss.

Mirroring protects against single points of failure in the storage array by mirroring the data to two (or more) disks. If one disk fails, the other continues to provide access to the data, while the first disk is replaced/repaired. When the disk is repaired, the mirror automatically restores the data and continues as before.

Protection from application or user errors

Highly available storage configurations do not protect data against software or user errors. Recovery from some errors typically requires the presence of adequate backups. It is imperative to maintain good backup procedures. The more frequent the backup, the faster the recovery in time of error.

Databases have special backup and recovery needs. Oracle databases in particular consist of many files that must be restored appropriately for a database recovery to take place. Although basic backup and recovery products can work for some Oracle databases, databases of a critical nature or significant size require Oracle-specific backup tools that can support online backups and recoveries, and can handle archived redo logs and other control files appropriately.

Whatever tool you use, taking frequent backups, testing recoveries, and cross-training DBAs are all crucial to effective Oracle administration.

Protection from disaster

If backup media are stored locally with the production systems, then localized disasters still threaten critical data. Protecting both the data and the physical backup media are essential elements of the broader disaster recovery plan. There are several approaches to protecting data from disasters, depending on how quickly you need to restart your business in the case of a disaster:

- **Off-site backups.** At the simplest level, maintaining off-site copies of backups gives you the ability to restart your critical data on a different server in case the primary site is inoperable.
- **Off-site replication.** If you need to restart operations very quickly after a failure, then a better solution is *replicating* critical data at an off-site location.

VERITAS solutions for data protection

VERITAS Software provides a complete, integrated family of products for protecting valuable data at all levels of the enterprise.

- **VERITAS Database Edition™ *for Oracle*®** protects data from disk failures through mirroring or RAID storage techniques, ensuring data availability even if a disk device fails. In the event of a disk failure in a redundant storage configuration, the hot relocation feature helps administrators replace disks without disrupting access.
- **VERITAS NetBackup™ DataCenter** is the industry's leading enterprise data protection solution for UNIX, Windows NT and NetWare. NetBackup provides fast, reliable, and data-center strength backup and recovery in environments that reach petabytes in size. NetBackup can protect the web servers, file servers, and application servers that drive your site.
- **VERITAS NetBackup™ *for Oracle*®** is the database-aware backup and recovery agent integrated with the NetBackup product family. NetBackup for Oracle can back up Oracle databases quickly while they are online and available. VERITAS NetBackup™ Advanced BLI Agent for Oracle only backs up data blocks changed since the last backup. Parallel backup and recovery operations enable high performance backups for very large databases.
- **VERITAS Volume Replicator™** leverages VERITAS Volume Manager™ to mirror data to remote locations over any IP network, reliably and efficiently for disaster recovery purposes. In the event of a primary site failure, the replicated site is both current and consistent.

"The Oracle/VERITAS combination means that we don't ever need to feel anxious about data integrity. Their powerful synergy lets us focus on what is really important: our site and our content."

— Sascha Linn, VP Networking and Operations, CheckOut.com

5. Protect the Applications

Site availability depends on the availability of many components. You must take an application-level view of availability and protect all of the components required to deliver online services to your customers.

The best approach for this is clustering software that leverages the hardware redundancy you built into your site in

Step 1 to provide sophisticated failover at an application level. There are different approaches to clustering, depending on whether or not the various nodes of the cluster will access the same data at the same time.

Shared nothing availability clusters

In a shared nothing cluster, each node of the cluster “owns” and accesses its own data, but has the potential to access other data on shared storage. In the event one node of the cluster fails, another node can then “take over” its applications.

This environment requires application-specific agents that monitor services on each node and handle failover policies.

Shared nothing clusters like this provide a high availability environment. In a web site, they can also provide a degree of horizontal scalability; you can have multiple web servers, for example, each accessing its own copy of the web site. This works well for the data that is read-only in the site, but introduces complexity if each node needs to update the same set of data.

Shared data clusters

Shared data clusters provide scalability as well as availability for a wide range of applications, and simplify administration. In a shared data cluster, multiple web servers can access the same copy of the web site; the cluster software arbitrates access. To implement shared data clustering, you need the software that arbitrates access and maintain cache consistency across the cluster.

Oracle Parallel Server (OPS) is one kind of shared data cluster specific to Oracle databases.

VERITAS clustering solutions

VERITAS Database Edition™/HA for Oracle®: The VERITAS Database Edition/HA for Oracle adds comprehensive application availability management for both hardware and software resources in clustered server configurations. Using **VERITAS Cluster Server™**, Database Edition/HA supports clusters of up to 32 servers in both Storage Area Network and traditional environments.

Using VERITAS Cluster Server, you identify the complete hardware/software resources required to serve an application, and design failover policies to handle the failure of any of these components. Cascading failover keeps your site available even if multiple servers or components should fail. Manual failover gives administrators the capability to rotate servers out of service to performance major administrative tasks, such as software upgrades.

VERITAS Global Cluster Manager™ monitors multiple, geographically distributed, Solaris-based VERITAS Cluster Server configurations, and manages fail-over of an entire cluster from one site to another anywhere in the world.

“HomeGrocer.com offers a 100% money-back guarantee for any customer that isn’t completely satisfied with our online grocery service. VERITAS Cluster Server has provided us with the ability to keep our website available, as well as, the scalability to effectively manage the enormous growth we’ve experienced over the past year – all of which helps us to stand behind the commitment to our customers.”

— Bob Duffy, Chief Information Officer

6. Plan for Change

Change is the one apparent constant in the Internet economy. Your site infrastructure must be able to adapt to change rapidly as you add new services or change strategies – providing fast time-to-market with reliable site performance.

Characteristics of an adaptable web site include:

- No reliance on a single hardware vendor for systems or storage.
- A multi-tiered architecture that enables you to add or modify services quickly and easily.
- Flexible administrative tools and policies that adapt to changing requirements.

VERITAS as a partner for change

Using the VERITAS suite of storage management products for e-businesses, you can keep your options open in terms of site infrastructure:

- Heterogeneous support. VERITAS products work with most major UNIX platforms and Windows NT/2000 platforms, so you can remain flexible in your choice of platform. They also work with a wide variety of storage devices, so you do not have to commit to high-end, proprietary storage solutions to get high end performance and flexibility. The VERITAS products help you manage the dynamic, multi-platform environment that characterizes many e-business sites.
- Storage Area Network (SAN) support. VERITAS Software has been among the first vendors to offer mature products for managing Storage Area Networks, with products for online SAN virtualization, heterogeneous clustering, and LAN-free backups. VERITAS is constantly delivering new products and technology to reduce the cost and complexity of managing Storage Area Networks, helping make the SAN architecture a viable reality in many installations.
- Industry leadership. VERITAS Software works closely with the major vendors in the e-business world, including Sun, Hewlett-Packard, Microsoft, and Oracle, as well as storage vendors, creating solutions that enhance those vendors' products and that drive the storage industry towards new levels of interoperability, performance, manageability and availability.

"Probably the most important partnerships today have been with Sun Microsystems and Oracle. Together with VERITAS, the three form the Mod Squad of Big Servers. Sun, Oracle and VERITAS are a common combination for customers looking for big box performance or building out muscle-bound Web sites for e-business."

— Chuck Phillips, Morgan Stanley Dean Whitter

Summary

The best way to jump start your e-business is to work with partners that understand the needs and realities of the unforgiving e-business environment.

VERITAS Software's long standing commitment to improving data availability in heterogeneous environments makes it the ideal partner for the growing e-business. VERITAS has built its business on optimizing and protecting critical data – offering “*Business Without Interruption™*” to the world's leading companies. Its integrated solutions can improve the performance, availability and manageability of almost any system/storage configuration.

In the online economy, VERITAS products are driving many leading e-business sites, because the performance, availability, and manageability they provide can make the difference between success and disappointment in e-business.



V
E
R
I
T
A
S

W
H
I
T
E

P
A
P
E
R

VERITAS Software Corporation
Corporate Headquarters
1600 Plymouth Street
Mountain View, CA 94043
650-335-8000 or 800-327-2232

For additional information about VERITAS,
its products, or the location of an office
near you, please call our corporate
headquarters or visit our website at
www.veritas.com