VERITAS™

BUSINESS WITHOUT INTERRUPTION™

(Up)Time is Money

The VERITAS High Availability Solution

Jonathan Martin

Product Line Manager, Clustering & Replication, VERITAS Software

The Need for Availability
Our Dependence on Access to Information

- **IT has become the business**
- **Downtime disrupts or cripples organizations**
- **Recovery from failure has to be faster every year**

12th June 1999: 22 hour outage
Cause: System Failure
Cost: $3M-5M + 26% drop in stock price

13th April 1998: 6-26 hour outage
Cause: Software Upgrade
Cost: $40M in rebates

24th February - 21st April 1999: Four 4 hour outages
Cause: Upgrade/Operator Error
Cost: $???M, $70M infrastructure investment

Costs of Information Loss

Lost Revenue Associated With Application Outages	
Application	**Cost Per Min**
Call location	$27,000
Number portability	$14,000
Enterprise resource planning	$13,000
Supply chain mgmt	$11,000
Electronic commerce	$10,000
Internet banking	$7,000
Universal phone services	$6,000
Customer service center	$3,700
POS/EFT	$3,500
Messaging	$1,000

IDC April '99

- **Cost of downtime is the easiest way to justify purchases**
- **Typical HA Customers:**
  - eCommerce/Web
  - Firewall
  - Financial Databases
  - Manufacturing
  - Telco
- **96.5% Uptime = 306.6 hours of downtime - Can your afford this?**

Information Availability Index
World-wide Availability Infrastructure

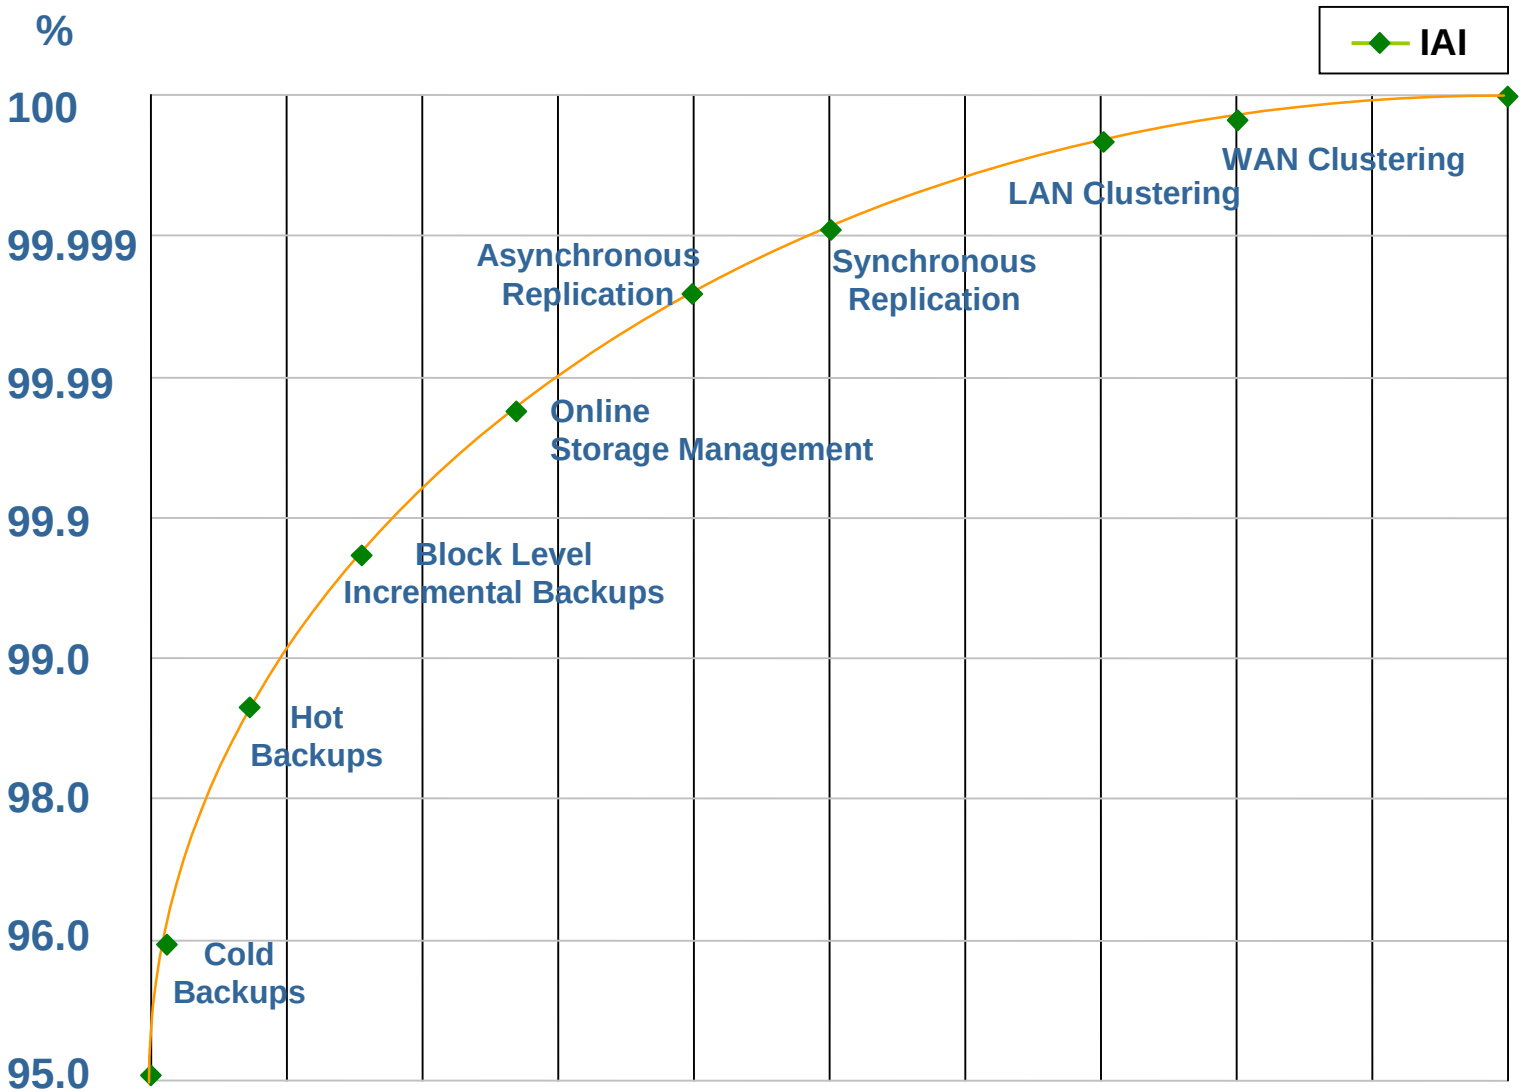

Information Availability Index

Windows NT 4.0 Availability

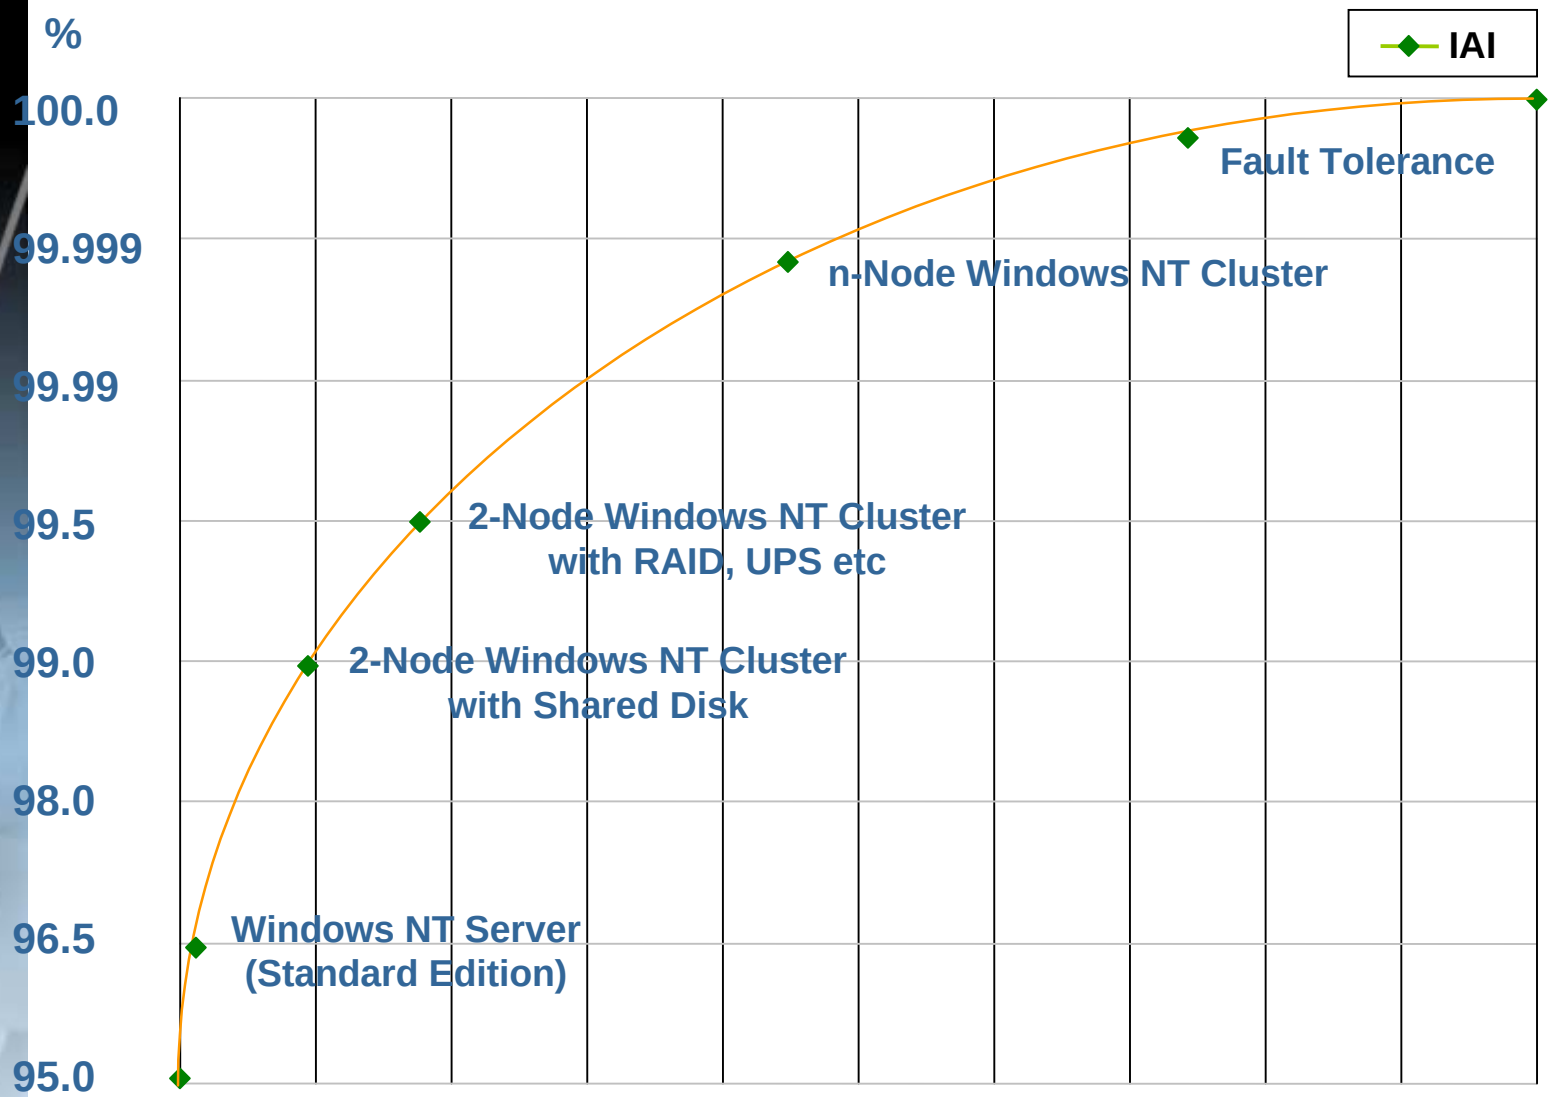

Source: Windows NT MSCS, Richard Lee, 2000

Recovery Time

Life after a hard stop

1	Safety Critical	99.9999%	45 seconds
2	Mission Critical	99.999%	
3	Business Critical	99.99%	
4	Task Critical	99.9%	
5	Task Non-Critical	Less than 99%	>87 hours

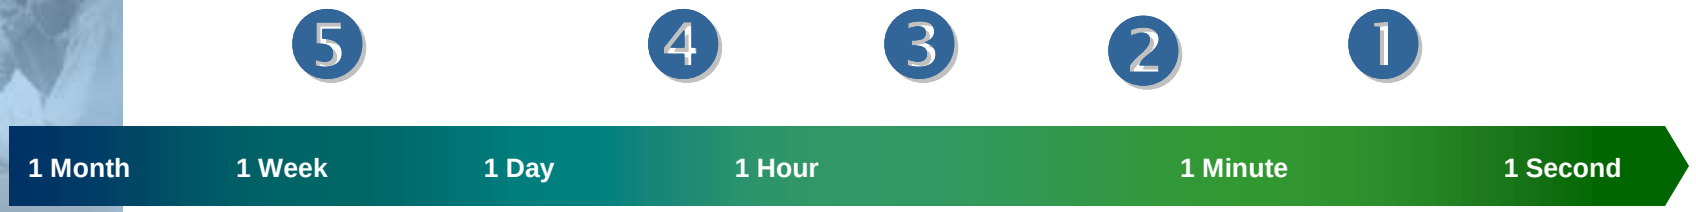

Recovery Time

Why Do Systems Fail?
The Causes of Downtime

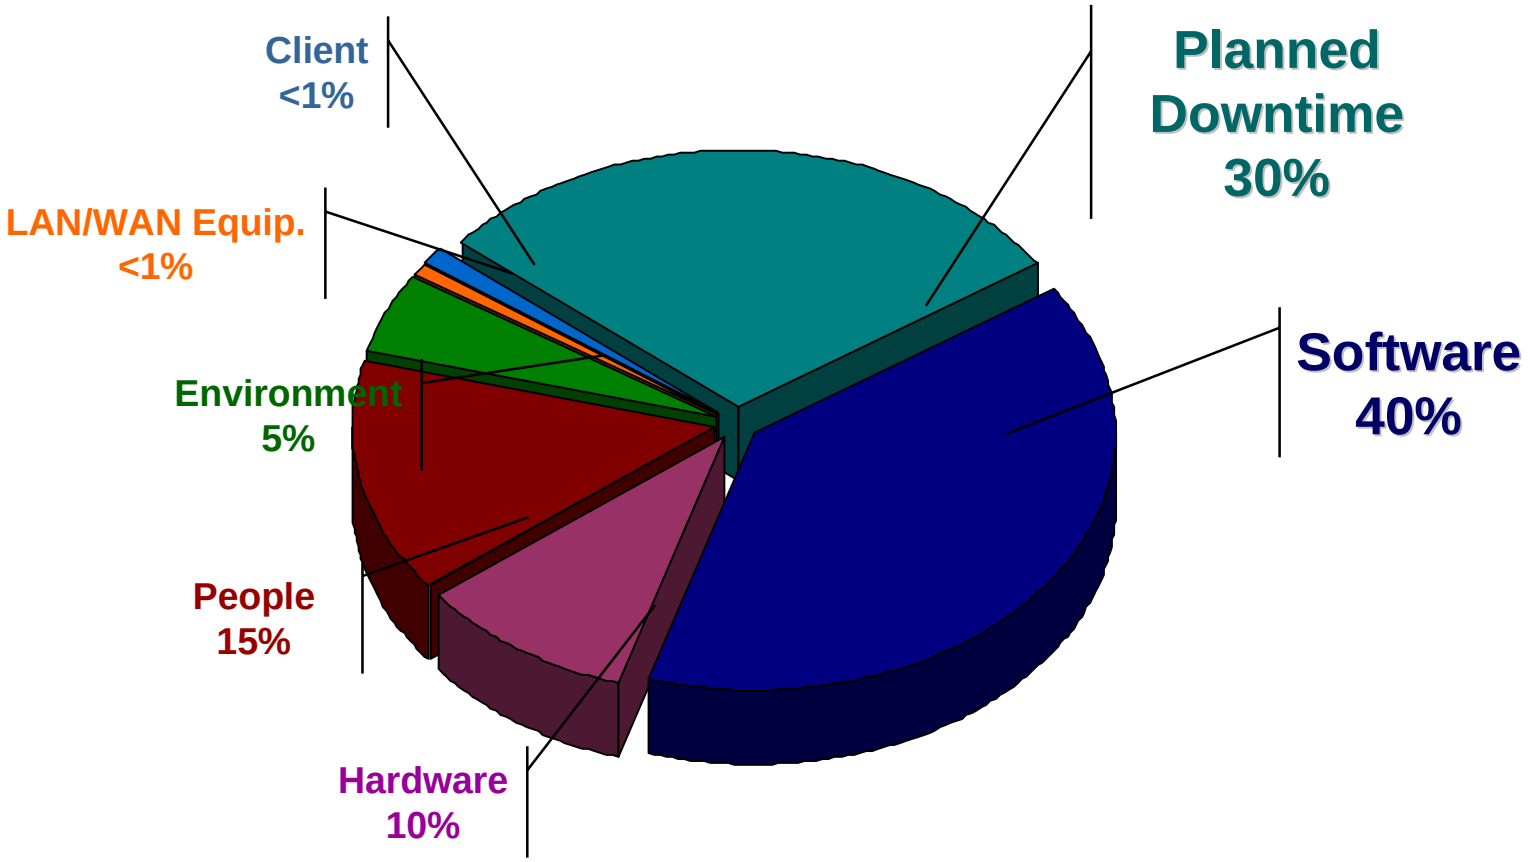

Source: IEEE Computer

Information Availability Index

Business Continuity Choices

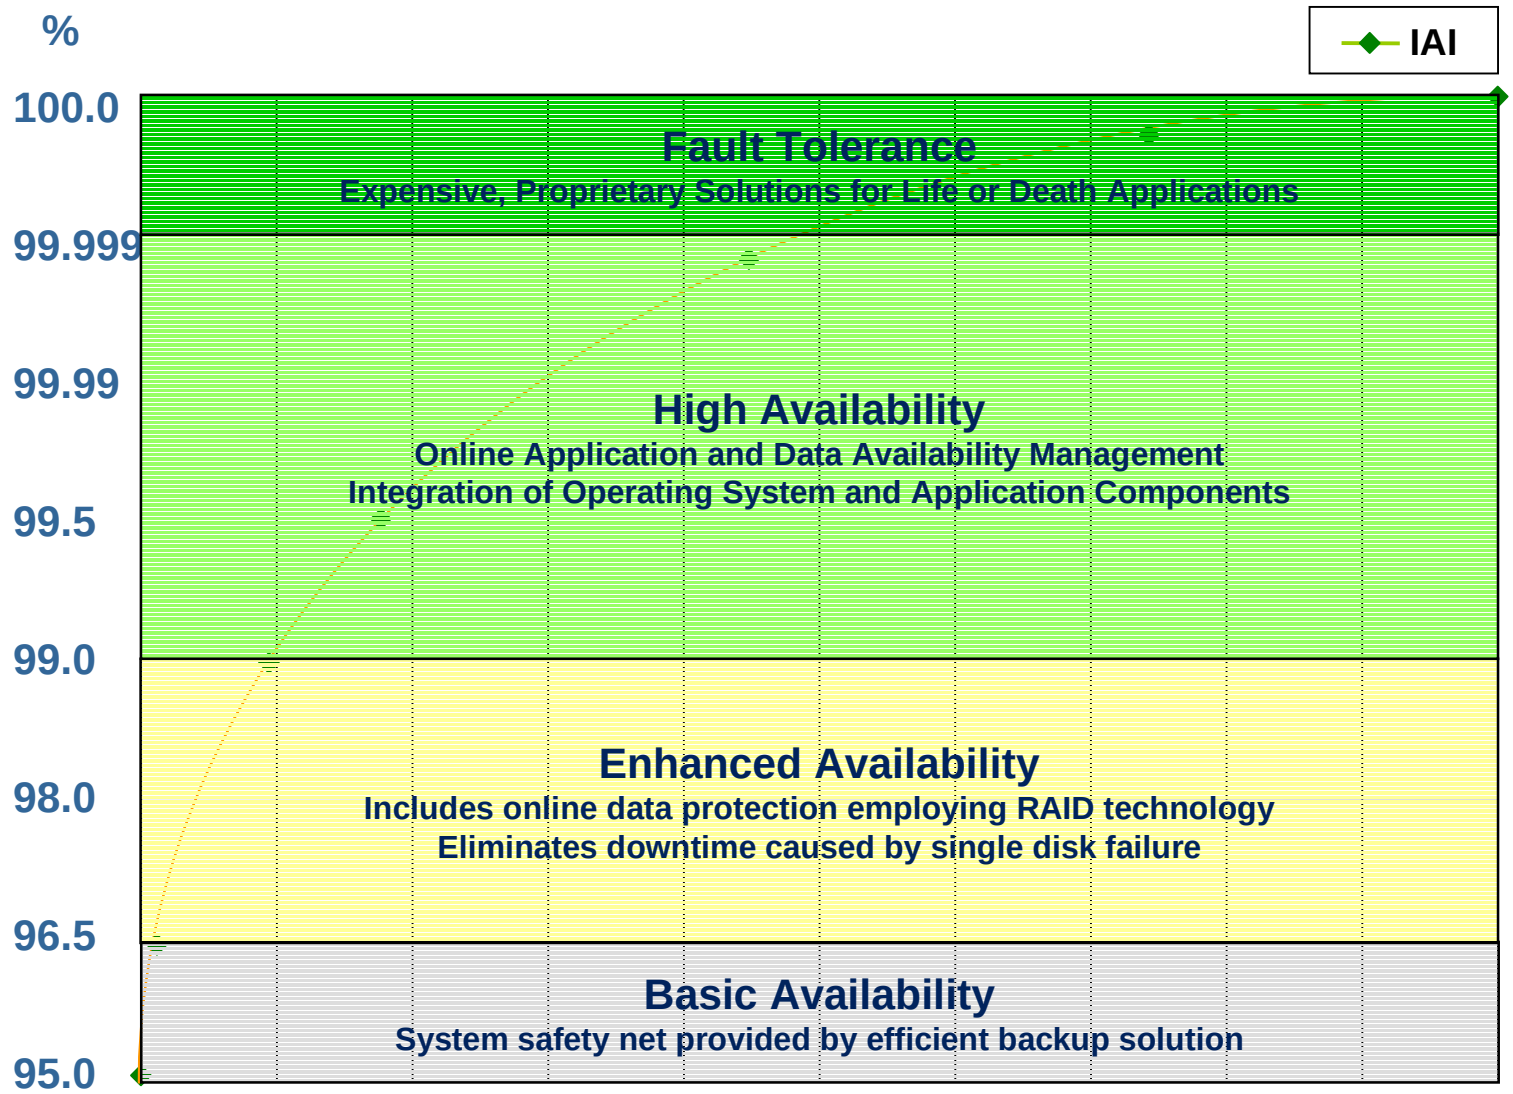

Information Availability Index

Business Continuity Choices

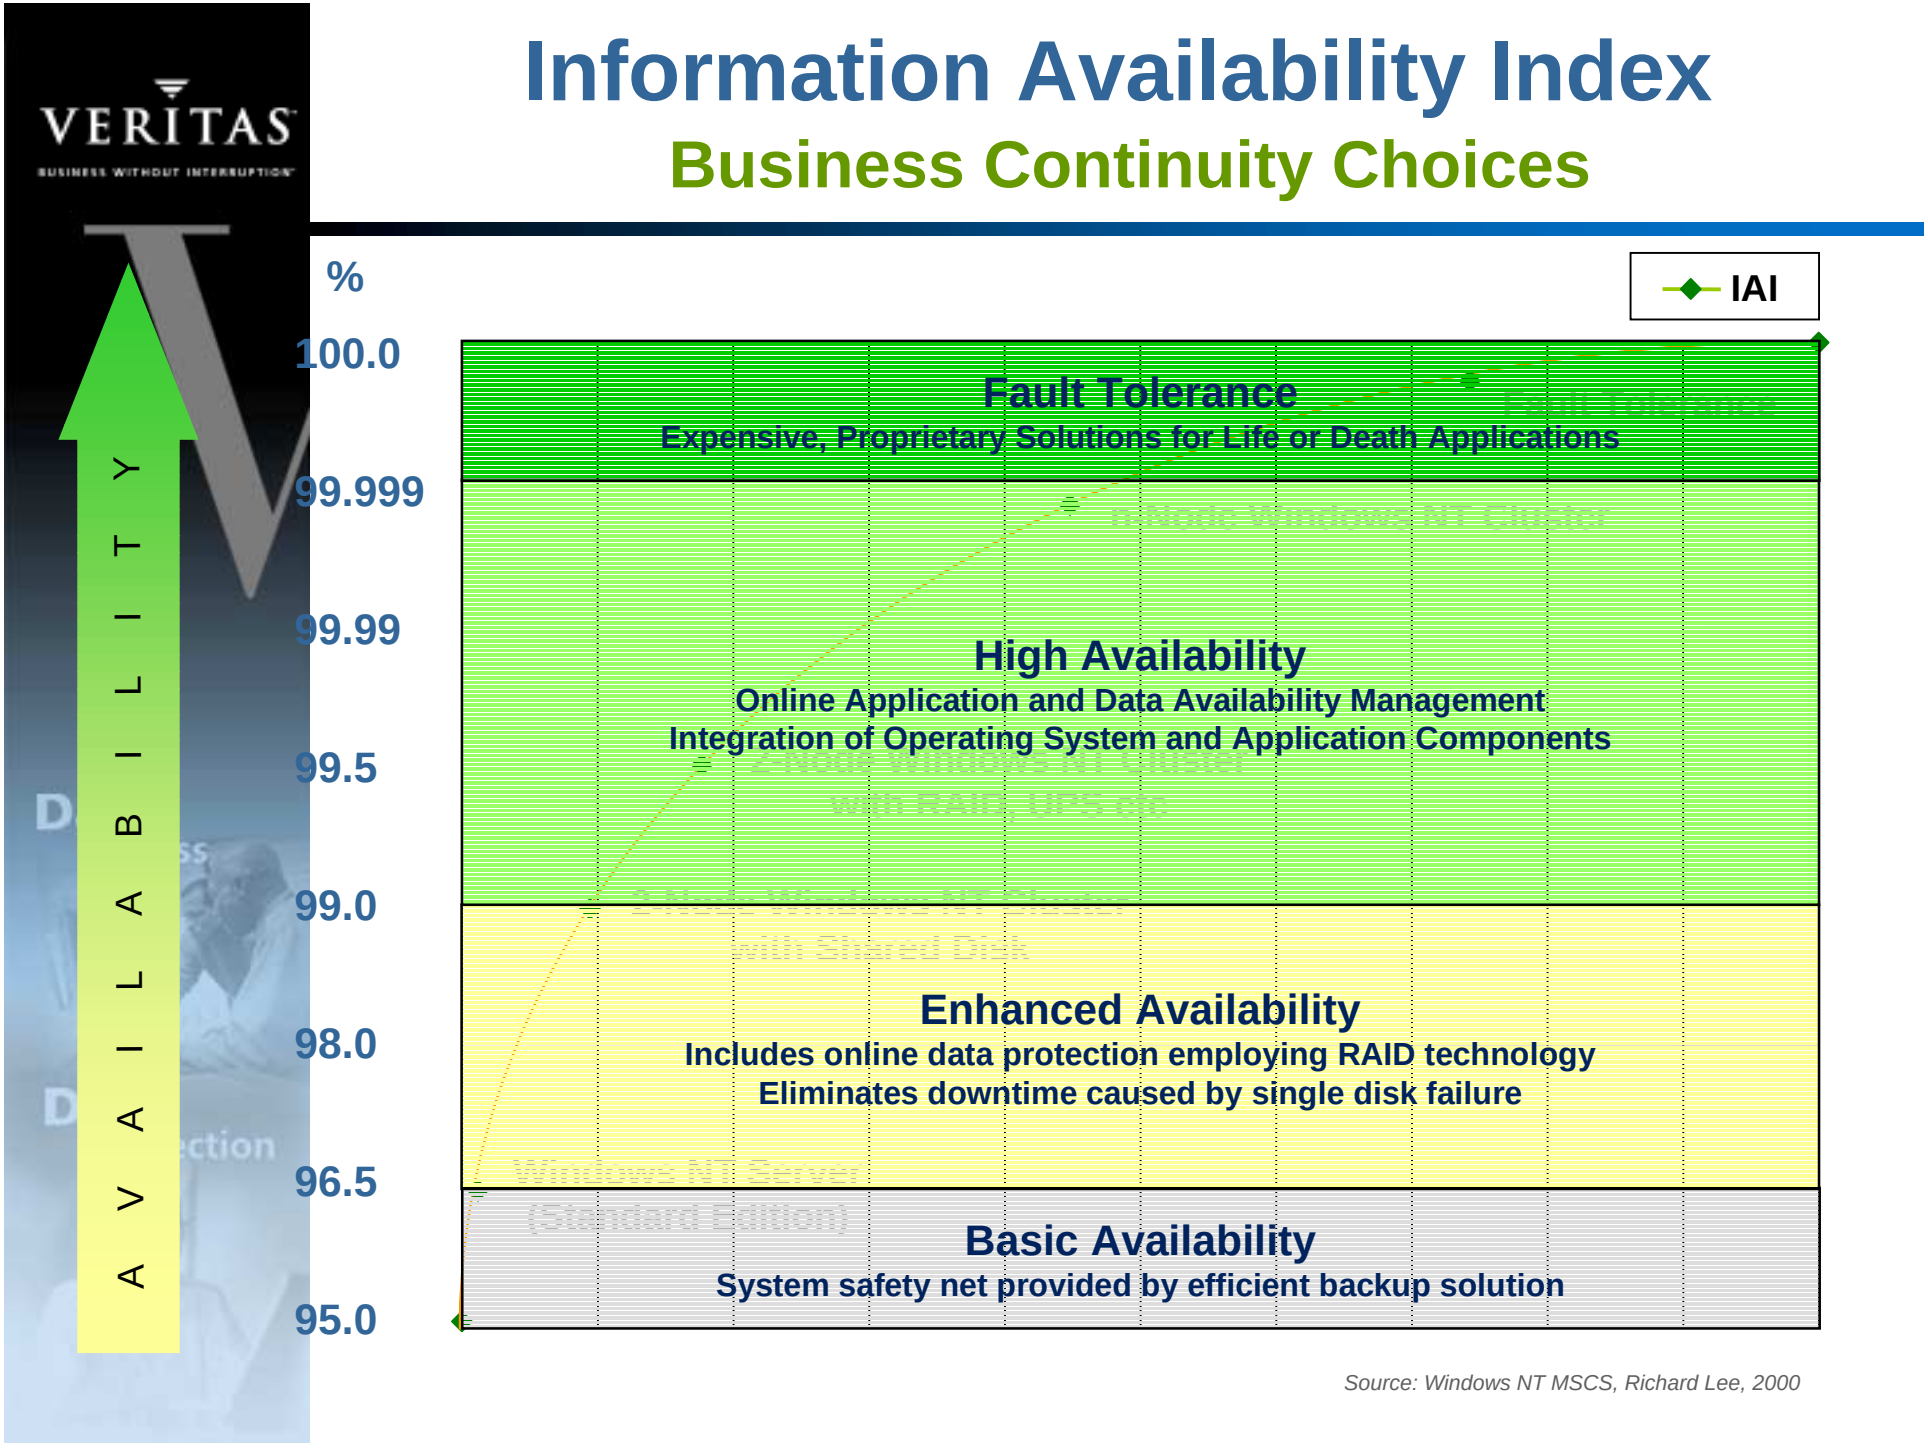

Source: Windows NT MSCS, Richard Lee, 2000

VERITAS High Availability

Platform Independent Information & Application Availability

- **5 components required to build a HA service :**
  - Redundant Hardware
  - HA System Software
  - Crash Tolerant Applications
  - Application Availability Management Software
  - Redundant Servers
- **VERITAS offer an extensive & integrated HA solution to help eliminate the majority of causes of downtime**

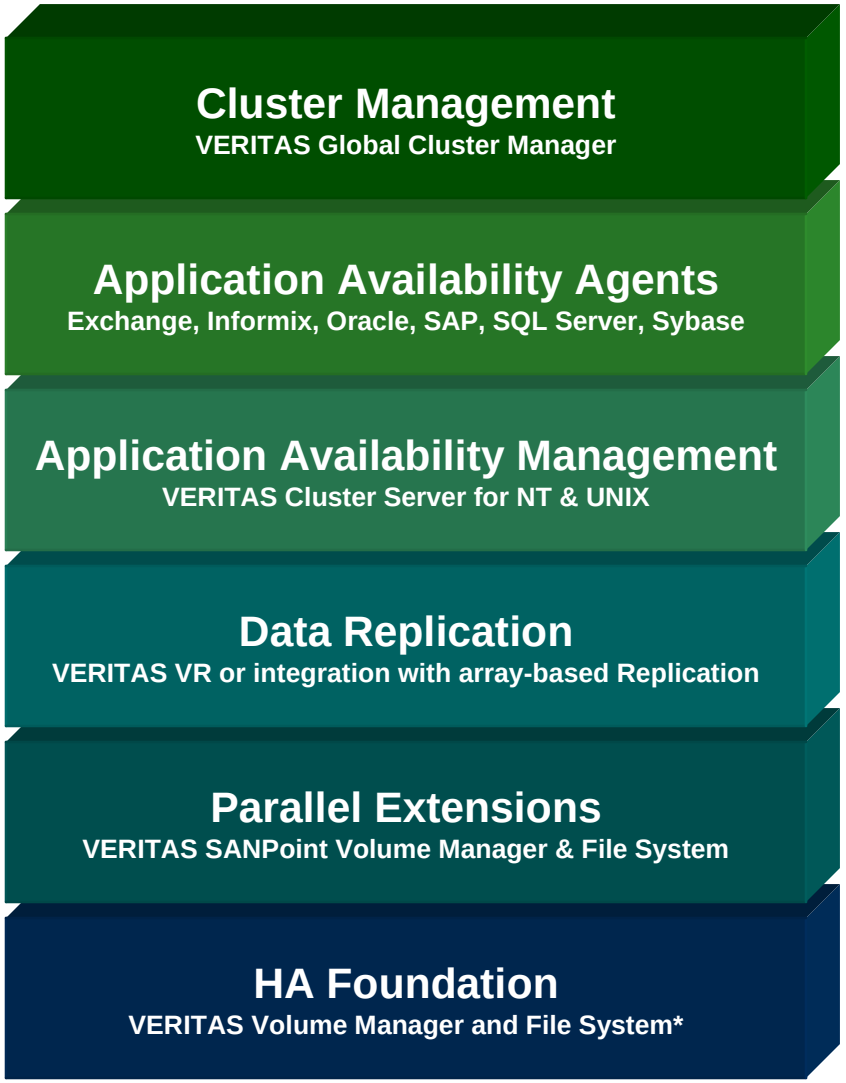

* UNIX only

Putting in a Firm Foundation

Removing the need for Planned Downtime

- **VERITAS Backup Exec & NetBackup**
  - Provides a safety net for your data
  - Fast application and information backup and recovery

- **VERITAS Volume Manager : Simplified HA Disk Management**
  - RAID-0, RAID-1, RAID-5, RAID 0+1 & RAID 1+0
  - On-line everything
  - Fast recovery
  - I/O optimization for file systems & databases
  - Online reconfiguration / re-layout
  - Optimise storage configuration depending upon access characteristics

Growing

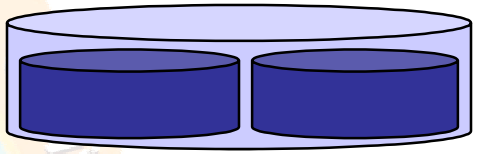

Online Relayout

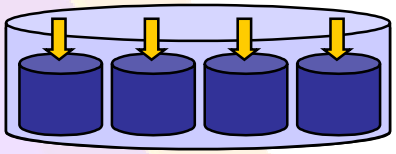

VERITAS SANPoint Volume Manager

Next Generation Clustered Volume Management

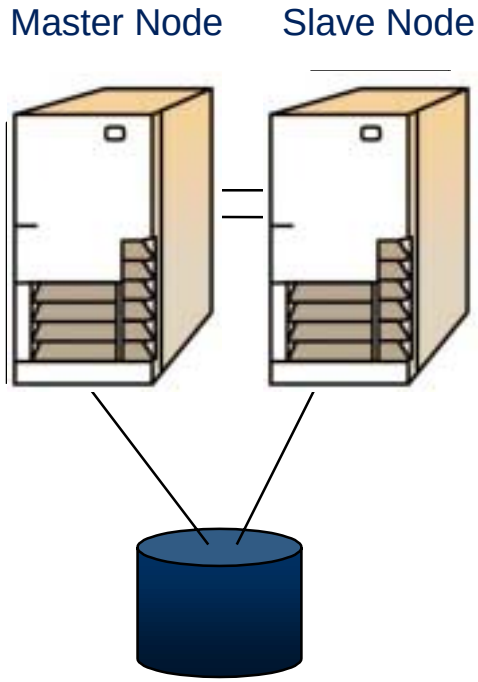

Cluster-shareable
disk group

- **Clustered version of VERITAS Volume Manager**
- **Simultaneous access to volumes from multiple hosts**
- **Managed from any host in the cluster**
- **Common logical device name**
- **Consistent logical view of volumes from any host**
- **Only raw device access supported from SANPoint Volume Manager**
- **Volumes remain accessible from other hosts after a single host failure**
- **Failover does not require volume migration**

VERITAS SANPoint File System

Advanced Clustered File System

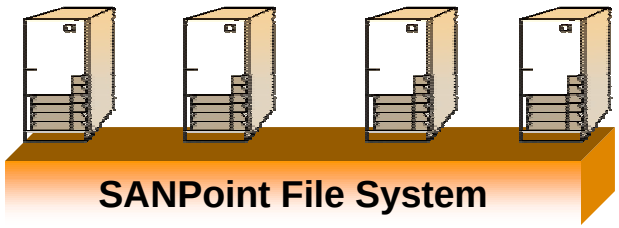

- Shared concurrent access to designated file systems
- Split client load across multi node resources
- Failover does not require FS restart

- Robust, high performance, general purpose parallel data service
- Shared concurrent access across multiple nodes
- Distributed lock management capabilities
- High availability features such as journaling and on-line mgmt
- Scales throughput as nodes are added
- Managed as a single image
- Part of SANPoint Foundation Suite/HA

High Availability: A Definition

- **Combination of two or more servers with appropriate middleware and interconnects to achieve some, if not all of the following:**
  - High levels of system reliability (availability)
  - Simplified administration of IT assets (manageability)
  - Increased processing power (scalability)

- **Other benefits may also be derived :**
  - Improved interoperability with other computing platforms
  - Reduced cost of ownership based on *server consolidation* and use of *virtual servers*
  - High return on investment and economic added value than standalone application servers

Failover: The Basis of HA

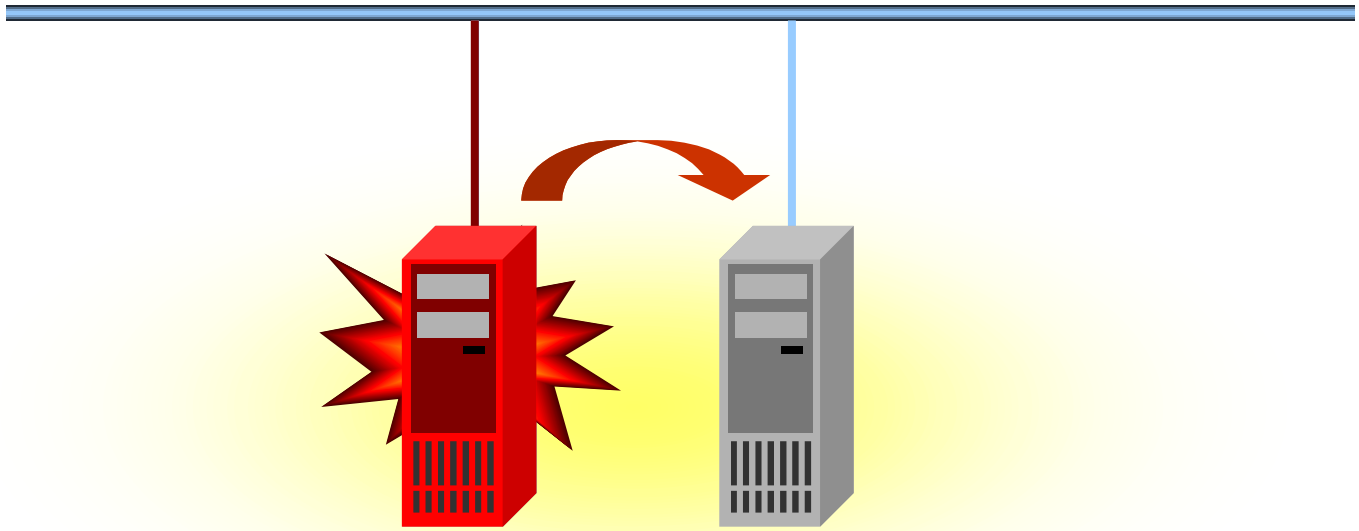

Application and Database (e.g. SAP with Oracle) and System Aware Failover

But 100% hardware redundancy required

High Availability Implementations

Traditional and Clustered HA

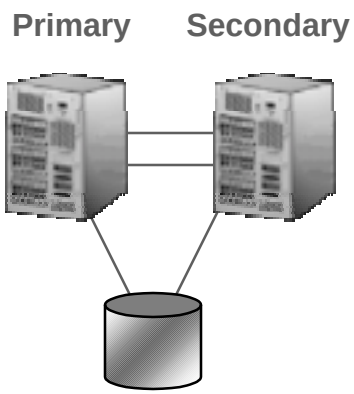

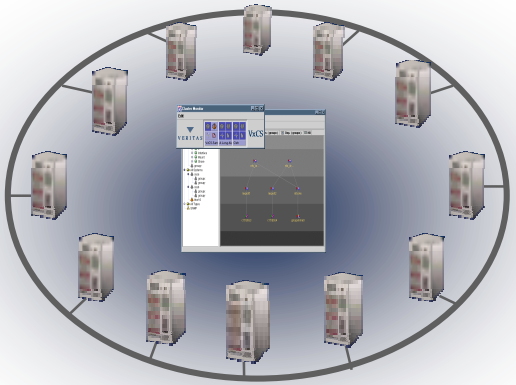

Traditional High Availability

- Primary hosts highly available application
- Secondary (“hot-standby”) only monitors until a failure occurs
- Limited Scalability
- Limited agent support for third party software and hardware (disk)
- No Wide Area HA Strategy

Clustered High Availability

- Second Generation HA technology
- Multiple servers, communicating together to ensure that clients maintain:
  - Accessibility to applications
  - Accessibility to storage / data
- Management from a central, cross platform interface

VERITAS Cluster Server
2nd Generation Clustered HA

- Sophisticated availability configuration management
- Clustered HA solution allowing 32 server scalability
- 1:1, Any:1 and 1:Any policy-based failover
- Application level cascading failover
- Fast Kernel-to-Kernel communication & failover
- Certified 3rd party support for SCSI, SAN & NAS storage
- Integrates with VERITAS Volume Manager & File System
- Centralized Cluster Manager console for HP-UX, Solaris and Windows NT

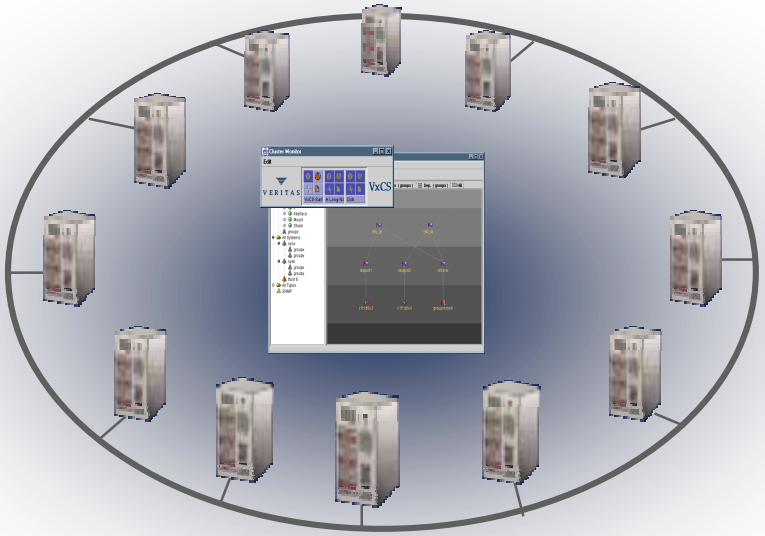

HEWLETT PACKARD Microsoft Sun microsystems

VERITAS Cluster Server
Application (Service) Level Failover

- **Monitored resources build a "virtual application service"**
  - application
  - IP address
  - disk or volume
  - server(s)
- **Policy based failover for load distribution**
- **Unlimited numbers of Service Groups**
- **Migrate one application without affecting others**
- **Cascading failover**
- **Build application dependencies**

A Service Group (SG) includes:

* One or more IP addresses
* One or more applications
* One or more disk groups
* A single physical host

VERITAS Cluster Server

The Application Agent "Glue"

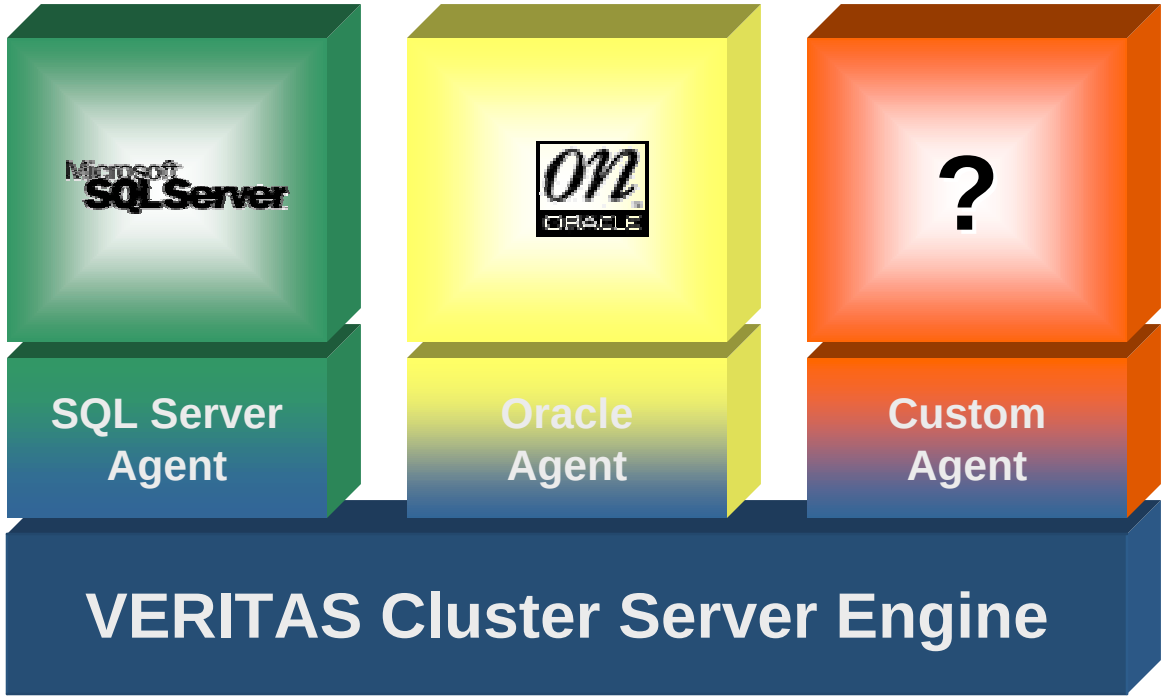

- **Application Agents**
  - Provide communication (translation) between Cluster Server Engine and the Application
  - Agents start, stop, restart and monitor applications

Application Integration:
Bundled Agents

- NFS v2 and v3
- CLARiiON ATF support
- vxdg (VM disk group) and VM
- Standard UNIX and NT files and processes
- File & Print Share
- Network (IP), haNIC
- Disk (raw devices)
- mount, share
- route (network routing)
- ServiceGroupHB
- SNMP
- Network Appliance

Pre-programmed agents speed deployment

Application Integration
Enterprise Agents

- BEA*
- Broadvision*
- Firewalls* (Checkpoint Firewall-1, Raptor*)
- IBM DB/2 EE/EEE*
- Informix ODS 7.x
- Lotus Domino Server 5.0
- Microsoft Exchange 5.5/2000**
- Microsoft-IIS 4.0/5.0**
- Microsoft SQL-Server 7.0/2000**
- Netscape SuiteSpot 3.5
- Oracle 7.x, 8.x, 8i
- SAP R/3 3.0 or later
- Solstice Internet Mail Server (SIMS)*
- Sybase 10.x, 11.x, 11.5
- VERITAS HSM*
- VERITAS NetBackup 3.4
- VERITAS Storage Replicator
- VERITAS Volume Replicator

* 2H2001
** Part of W2K release

Now and Beyond:
"Storage Area Network" Architecture

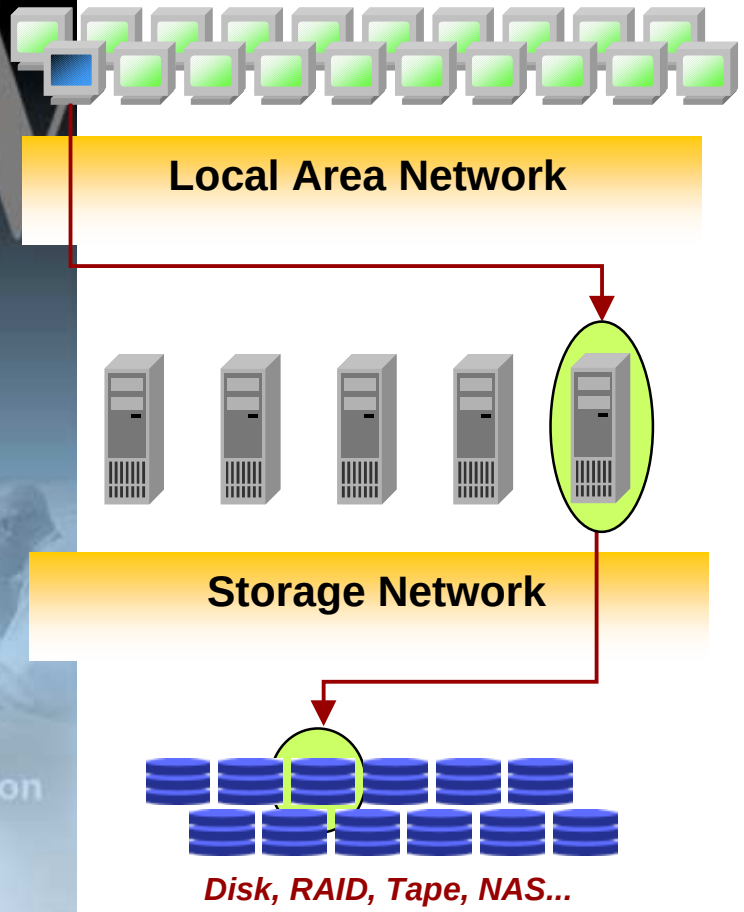

- "Any-to-any" storage connectivity
- Ubiquitous data access
- Higher availability — at significantly lower cost
- Easy and inexpensive scalability
- Reduced management costs

SAN Support

An infrastructure for scalable clusters

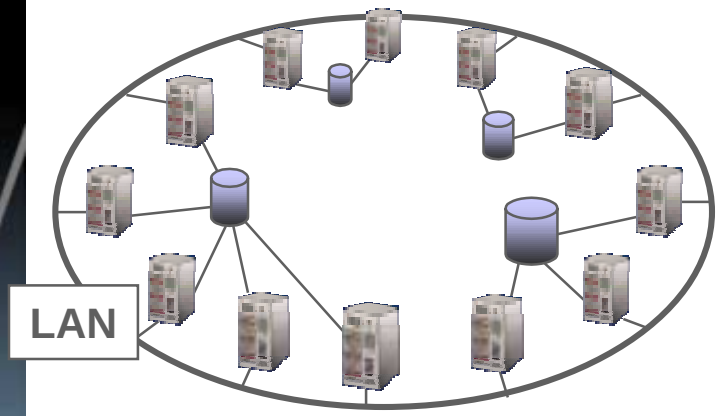

Shared Disk Model

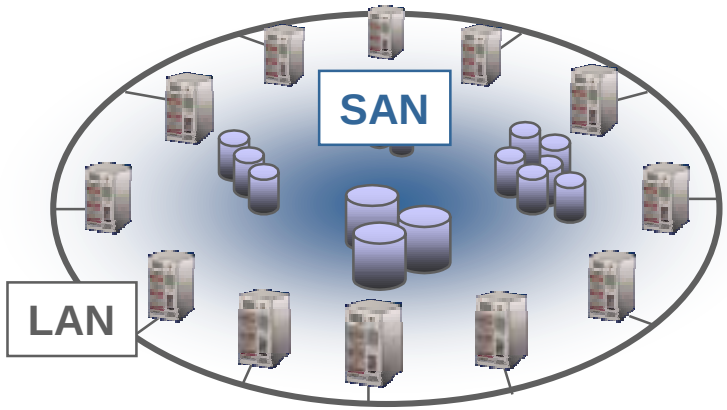

SAN Model

- SAN forms an infrastructure for scalable Clustered HA
- SCSI disk reservation agent prevents data corruption on shared SCSI disk due to network partition
- VERITAS Volume Manager Disk Group Import/Deport integration
- VERITAS iLAB deliver referencable SAN configurations

Cluster Management

- Single console for cross platform VERITAS Cluster Server (VCS) administration
- Manual failover for administrator initiated workload balancing
- Secure *roles*-based administration
- Wizard driven for ease of use
- VCS Command Center for point and click Service Group construction
- Administration traceability using VCS Log Desk
- Tunable logging and debugging
- SNMP alerting for integration with frameworks
- Local or Remote Administration

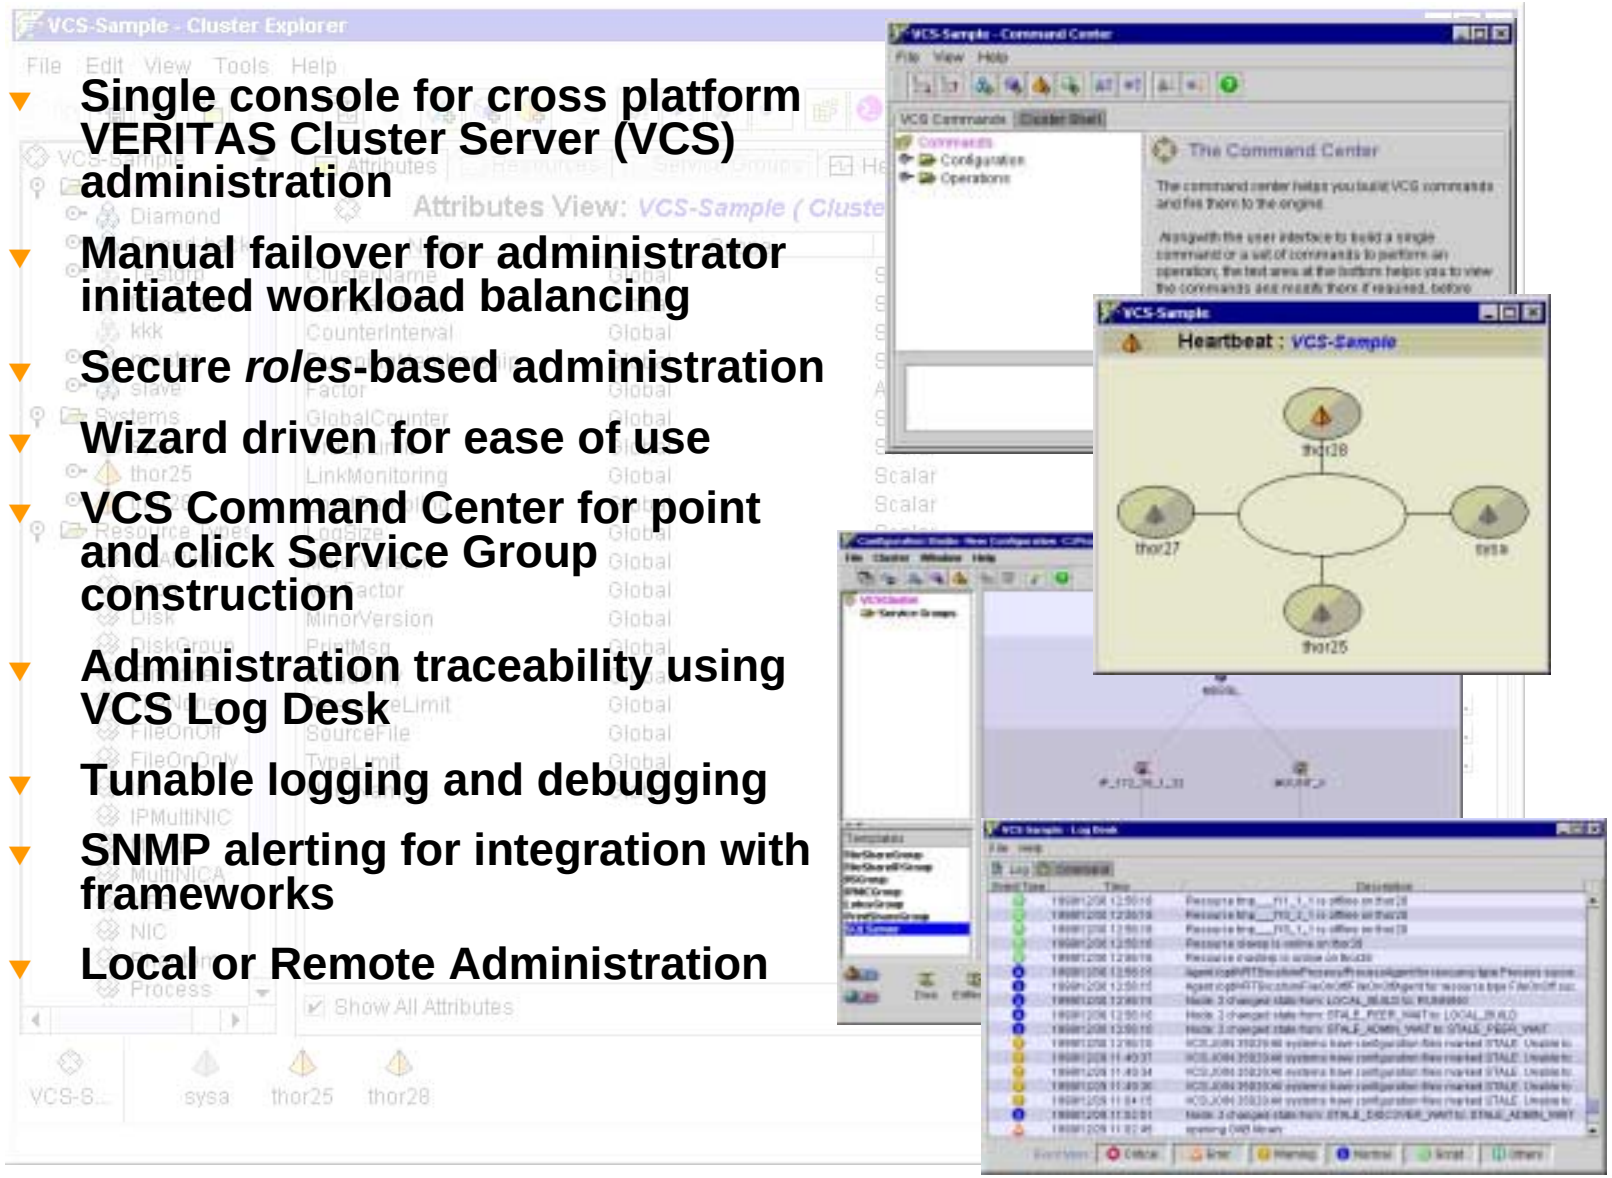

VERITAS™

BUSINESS WITHOUT INTERRUPTION™

Wide Area HA Solutions

Wide Area High Availability

Integration of Management, Clustering and Replication

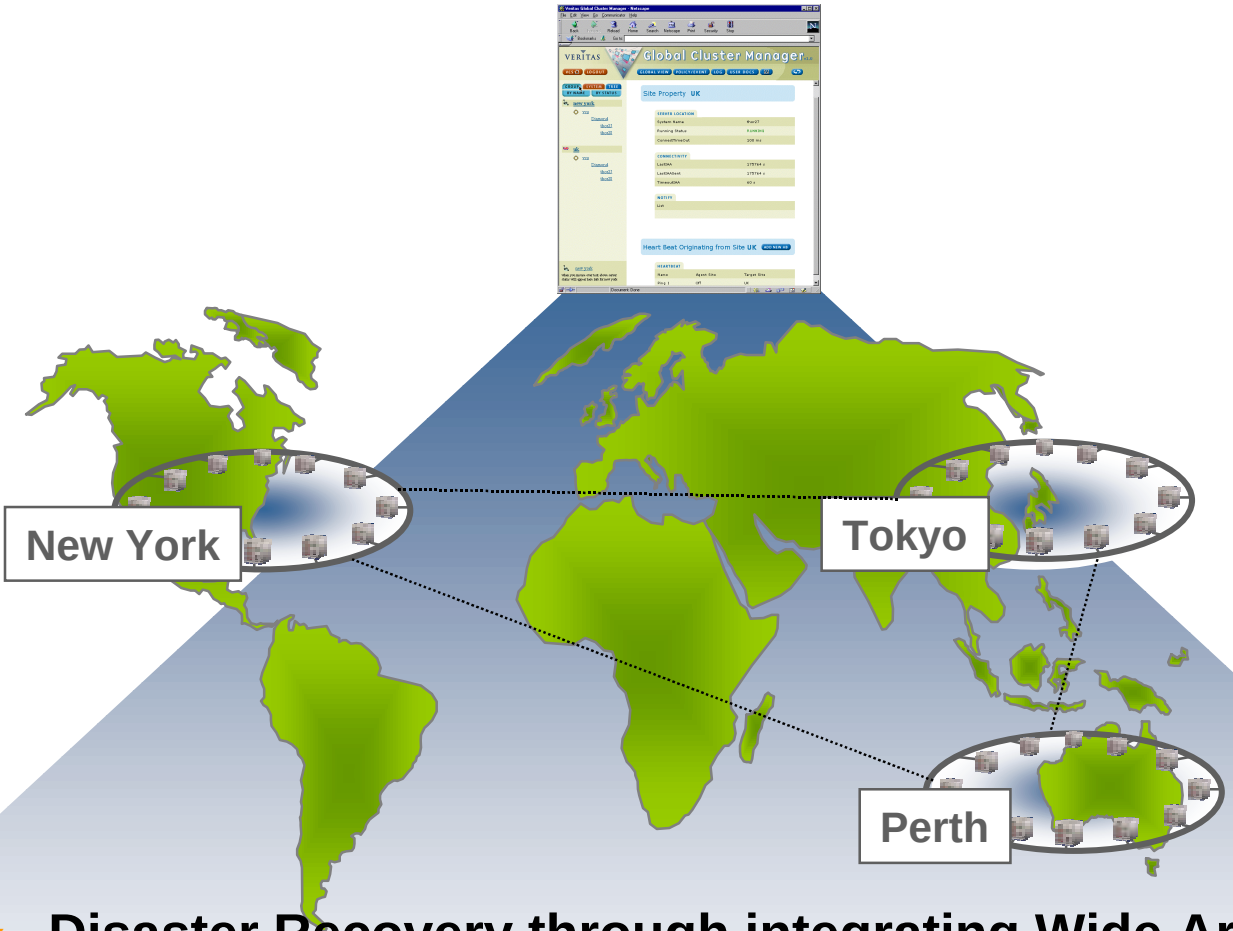

- **Disaster Recovery through integrating Wide Area Failover with Data Replication**
- **Provides a single point of administration and management for multiple geographically distributed clusters**

VERITAS Replication Family

Disaster Recovery/Offhost processing

New for NT/W2K in 3Q2001

- VERITAS Volume Replicator
  - Real-time disaster recovery on DBMS or FS
  - R/W at primary only
  - Off-host processing via mirror break off at secondary
  - Modes: Full synchronous, dynamic synchronous, asynchronous

Distributed Hosting

- VERITAS File Replicator
  - Flat file content or software distribution (no DBMS)
  - R/W everywhere
  - Modes: Synchronous (future: Sync, Async, periodic, or event driven)

- VERITAS Storage Replicator
  - Content Distribution and Disaster Protection for Windows NT environments
  - R/W primary, read-only target(s)
  - Modes: Real-time asynchronous, scheduled or periodic

Volume Replicator - Flexibility

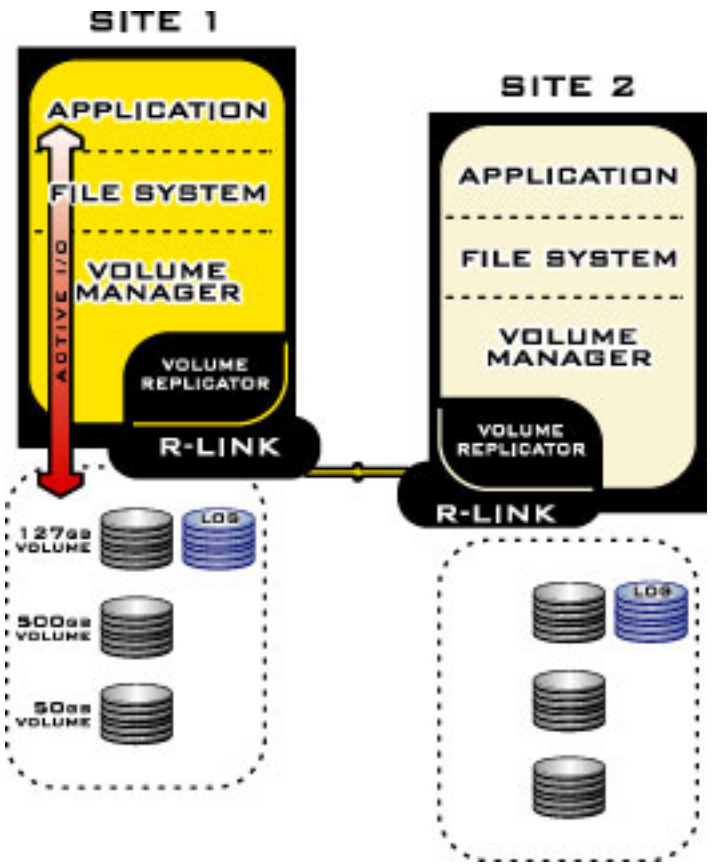

- Maintain ALL VxVM Online Management
- Supports ANY DBMS or FS in both Sync and Async Modes
- Supports ANY Storage (EMC, HDS, Sun, MTI, STK, IBM…)
- NO Distance Limitations
- Any Storage Layout (w/ VxVM)
- Async override option
- Initialization options
- Can replicate volumes that span storage arrays
- Shared network support
- Scales to 32 locations

Global Cluster Manager

- **Global Cluster Manager offers an enhanced HA feature set :**
  - Cluster Scalability
    - GCM provides a framework for distributed applications
    - Monitors application services, and can be configured to fail over faulted instances to alternate site
  - Consolidated Monitoring and Management
    - Manage multiple clusters of different architectures (e.g. HP-UX, Solaris and Windows NT) from a single console.
    - Co-ordination of events and management or responses
  - Administrative Migration

    New for NT/W2K in 3Q2001

    Migrate services from one cluster to another to remove consequences of planned downtime or the need for “follow-the-sun” applications e.g. support
  - Disaster Recovery

    New for NT/W2K in 3Q2001

    Physically remote clusters provide continuous availability of application services and data, should a site become unusable

Global Cluster Manager

Integration of Management, Clustering and Replication

Two geographically dispersed sites

Linked by public carrier wide area network

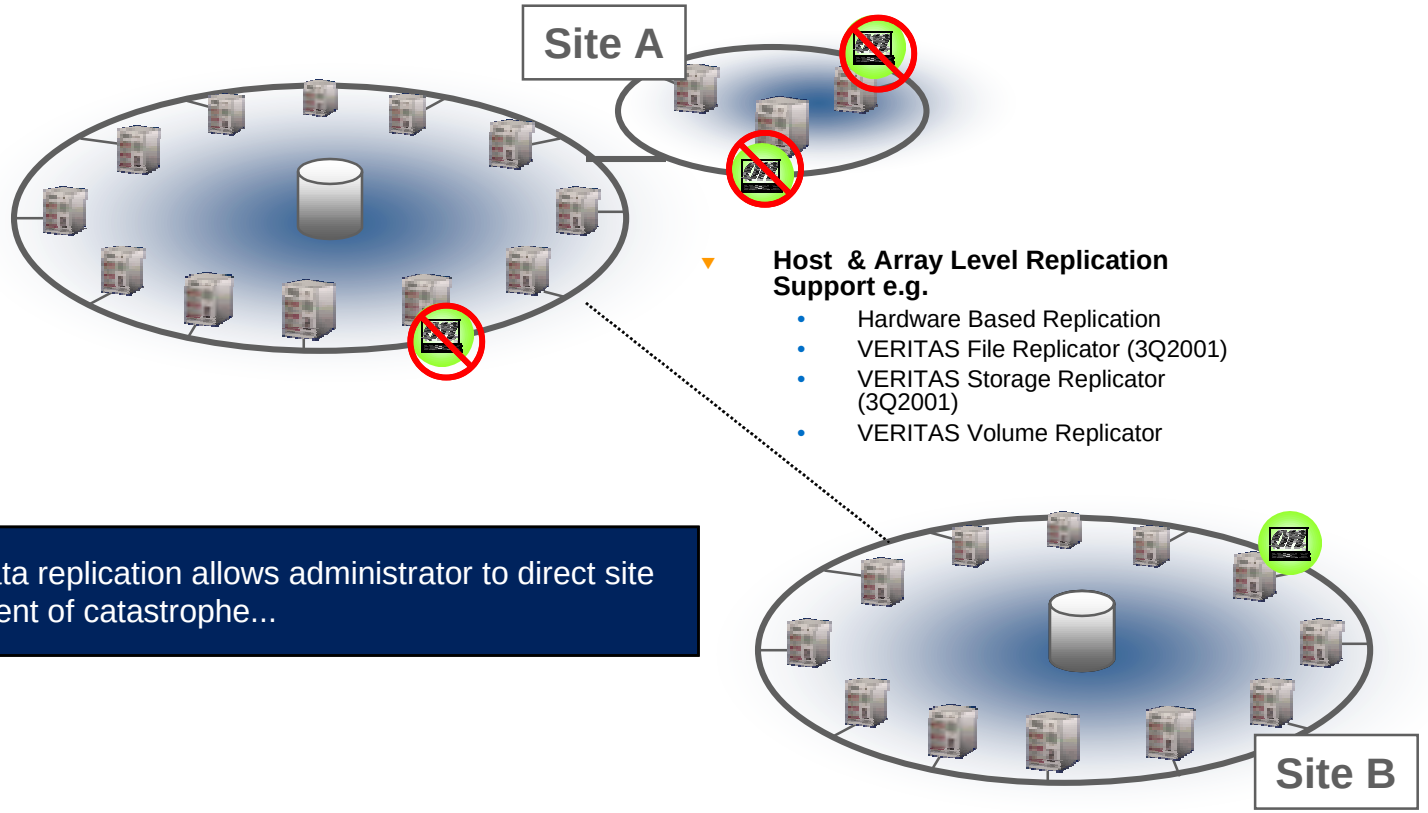

- **Host & Array Level Replication Support e.g.**
  - Hardware Based Replication
  - VERITAS File Replicator (3Q2001)
  - VERITAS Storage Replicator (3Q2001)
  - VERITAS Volume Replicator

Integration with data replication allows administrator to direct site level failover in event of catastrophe...

Global Cluster Manager

Architecture Overview

Cluster Slave

- A Cluster Slave process resides on each cluster
- The Cluster Slave extracts cluster, system, group config., and state information from VCS and forwards it
- The Cluster Slave can be configured as a group for high availability.

Site B
Three Clusters

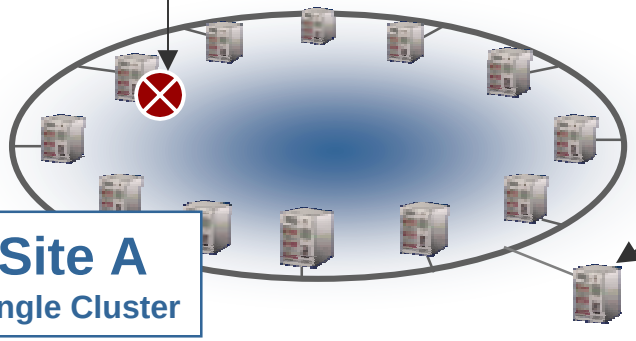

Site A
Single Cluster

Site Master

- One Site Master per site
- Standalone or as part of a cluster, but is visible to other GCM sites
- Consolidates information and shares it with other Site Masters
- Each site master is aware of the configuration and state of all of the sites to which it is connected

Global Cluster Manager

Architecture Overview

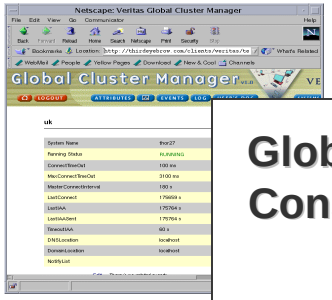

Global Cluster Manager Console

A graphical user interface, similar to the VCS GUI, where you can view the configuration and states of all sites, clusters, and groups being managed by GCM. You can use the GUI to track events and issue directives such as bringing groups online and migration

Site B
Three Clusters

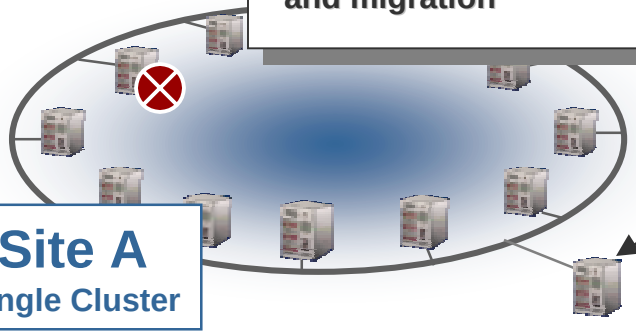

Site A
Single Cluster

Site Master

- One Site Master per site
- Standalone or as part of a cluster, but is visible to other GCM sites
- Consolidates information and shares it with other Site Masters
- Each site master is aware of the configuration and state of all of the sites to which it is connected

Summary

- HA solution that scales with your requirements
- Central management of multiple clusters
- Application level availability management for uninterrupted access to data, applications and other network services
- Eliminates the majority of causes of downtime by integrating with products like VERITAS Volume Manager and VERITAS File System
- Cross-platform NT and UNIX solution for reduced ownership cost
- Backed up with enterprise service and support from the industry's leading vendor