




Fibre Channel Clustering

“The Next Generation”




Unisys Interface Requirements

Independent Studies

System Platform Group

Enterprise Servers



- Customer Focus Group
- Industry Standards Organization
- OEM Strategies
- Consultant Analysis

Storage
Clustering
Interconnect

Storage Systems Group


Internal Market

OEM Market




Host to Device
Host to Controller
Controller to Device


Unisys
PrimeStor



Leverage Investment Multi Functional Standard




Storage




Clustering - Definition

- A cluster is a group of computers and storage devices that function as a single system
- A cluster enables multiple independent servers to be interconnected so that any server can perform any task
- A cluster can be complementary or an alternative to SMP technology





Clustering - Value Proposition

- Mission critical computing transitions from the “*glass house*” to a “*distributed computing model*”
- Clustering provides three fundamental advantages
 - High Availability
 - No single point of failure stops the system
 - Modular Scalability
 - System upgrades with investment protection to address growth opportunities
 - Ease of Management
 - Operations management centralized




Clustering - Architecture





Clustering - Marketplace Survey

- Capabilities required for cluster deployment:
 - Fault tolerant (Range of high availability options)
 - Guaranteed message delivery
 - Low latency (Deterministic option)
 - High bandwidth
 - Broadcast and multicast (For large cluster configurations)
 - Scalable growth (Support for 128+ processor configs)
 - Dynamic reconfigurations (Without interruption)
 - Load balancing and cluster channel prioritization options
 - Management status reporting



Clustering - Capability Matrix

<i>FEATURE</i>	<i>DESCRIPTION</i>	<i>EXAMPLE</i>
High Availability	Enhanced application and service availability	Unisys Network Attached Storage
Modular Scalability	Partitionable applications can grow beyond the capacity of any single server	Oracle Parallel Server
System Balancing	Makes application performance independent of execution site	Digital VMS Cluster
Single System Image	Unified management, user, and application view across multiple servers	IBM Parallel Sysplex




Unisys Network Attached Storage


Fibre Channel Clustering

- NT Mirrored Cache
- Heartbeat Check
- System Failover

Fibre Channel Storage


- Dynamic RAID






Clustering - Implementations

- LAN/WAN deployment to connect systems into a cluster
 - Can't address large clusters and high performance applications
 - Can't provide low latency and high bandwidth
 - Can't provide robust feature content




Clustering - In the Marketplace

- Digital Clusters
- Tandem ServerNet
- Sequent NUMA(Non-Uniform Memory Access) Servers
- Vinca Standby Server
- Microsoft Wolfpack
- Intel VIA (Virtual Interface Architecture)
- Intel Architecture-based SHV
 - (Standard High Volume) Server Platforms
- Unisys Network Attached Storage



Clustering - In the Marketplace

- Clustered Databases
 - Oracle, Informix SQL Server
- Clustered Transaction Processing
 - BEA Tuxedo, Tandem ServerWare, Microsoft MTS
- Clustered Shared Memory
 - ccNUMA, Digital Memory Channel



Clustering - Software




- Microsoft Wolfpack


- Phase 1 supports 2 node failover

- Hardware or software failure triggers the node switch

- Phase 2 will support multinode failover

- Higher performance and availability






Clustering - VIA Initiative

- Intel
- Microsoft
- Compaq


Mission critical high performance clustering
alternative at a lower TCO than:

- Proprietary Mainframe
- RISC based Cluster




Clustering - Implementations

- Virtual Interface Architecture (VIA)
- Who:
 - Led by Intel, Microsoft, Compaq, and 50+ companies
- Why:
 - Need for an open standard architecture for cluster computing that utilizes a lightweight protocol to overcome software overhead in the legacy protocol stacks
- How:
 - Achieved by avoiding intermediate copies of data and bypassing the OS when sending & receiving messages



Clustering - Implementations

- Virtual Interface Architecture (VIA)
- What:
 - VIA defines mechanisms for low-latency, high bandwidth message-passing between interconnected nodes and interconnected storage devices
 - The VIA spec provides standard hardware and software interfaces for cluster communications
 - The VIA spec is media, processor, and OS independent
 - The hardware interface will be compatible with standard networks such as ATM, Ethernet, and Fibre Channel




VIA - Goals and Objectives*

- Fulfill customer requirements by providing reliable, cost effective building block products
- Ensure the adoption of standards-based clustering technology
- Accelerate the expansion of the clustering market


* Compaq

Unisys
PrimeStor



Clustering - Implementations


- Tandem ServerNet
 - VIA standard compliant
 - 2 node Wolfpack support
 - Multinode distributed database support
 - SCSI storage environment (or replacement)
 - 64 to 1024 byte packets
 - 125MB/sec links
 - 64 bit addressing



Clustering - Implementations


- Tandem ServerNet Quote on Fibre Channel

“Fibre Channel has been optimized for long transfers and has unacceptably long latencies for short IPC messages. Fibre Channel also lacks the ability to pull data with read operations, making them impossible to use for efficient remote PCI bridging.”





Clustering - Directions

- Evolve from the current 2 node application failover solution to multinode parallel application clusters
- More robust messaging techniques between cluster nodes
- New generation of enterprise solutions will be enabled by efficient lightweight protocol stacks




Clustering - In the Marketplace






Clustering - In the Marketplace

- Clustering “Killer Apps”
 - Data Warehousing
 - Decision Support
 - Transaction Processing
 - Lotus Notes
 - E-Mail



VIA Clustering - Applications

- 1998 Software Product Development utilizing the VIA specification
 - SAP
 - Baun
 - PeopleSoft



Amdahl Cluster - Application Profile


1997 Informix User Group Conference

Hardware:

Amdahl NT Servers (4 node cluster)

Software:

Informix Online XPS (Extended Parallel Server)



Unisys Cluster - Customer Profile


Customer Benchmark

Hardware:

16 Unisys NT Servers (Each Server is a 10 Processor SMP)

Software:

Oracle Parallel Server and Tandem ServerNet



Clustering Directions - Oracle


Oracle 8 utilizing VIA for clustering

Hardware:

***Intel 64 bit Merced Compaq, Hitachi, HP, Sequent, Seimens,
Unisys + Others***


Software:

Support for NT and Unixware




Clustering Summary

- Open standards for clustering will emerge in 1998
- Open systems application and database development will be architected to take advantage of clustering
- The major “industry players” are committed to accelerating the adaption of cluster computing



Clustering - Possibilities

- Do we create a FCA Cluster SWIG?
- Do we engage T11 support for focused support of cluster standards?
- Do we develop a “core team” of companies to develop a Fibre Channel cluster strategy?
- Does FCA pursue tradeshow and industry events involved with clustering technology?



Clustering - Related Documents

1. U-Net:

"A User level Network Interface for Parallel and Distributed Computing"

Thorsten von Eicken, Anindya Basu, Vineet Buch, and Werner Vogles
Department of Computer Science. Cornell University

2. "NOW" (Network Of Workstations)

David E. Culler

Computer Science Department at University of California Berkeley.


3. Software Support for Virtual Memory Mapped Communication

Kai Li, Edward Felten, Liviu Iftode and Cezary Dubnicki
Department of Computer Science, Princeton University

4. The Paragon Message Passing Interface Paper

Paul Pierce & Greg Regnier

SHPCC94



Clustering - Related Documents

5. "In Search of Clusters"

Greg Pfister. Prentice Hall

6. "An Analysis of TCP Processing Overhead,"

D.D. Clark, V. Jacobson, J. Romkey, and H. Salwen
IEEE Communications Magazine

7. "The Importance of Non-Data Touching Processing Overheads in TCP/IP"

J. Kay and K. Pasquale

Technical Report, VI Architecture Diego

Department of Computer Science and Engineering, University of California

8. "LogP Quantified" The Case for Low-Overhead Local Area Networks

K. Keeton, T.E. Anderson and D.A. Patterson

Hot Interconnect III