CS514: Intermediate Course in Computer Systems

Lecture 8: Sept. 24, 2003

Cluster Computing

Guest lecture from Werner Vogels

Today: Cluster Computing

Werner Vogels
Dept. of Computer Science
Cornell University

Agenda

CS514

- History of Cluster Computing
- Major Cluster Systems
- Technical Challenges
- Software Architectures
- Cluster Management Systems
  - MSCS
  - Galaxy
- Student Research Projects in Cluster Computing

What do I want you to know?

CS514

- Why and how clusters are used?
- What is the difference between parallel and enterprise cluster computing?
- What are the major issues in hardware and software?
- What is a Cluster Management System?
- What do I need to do to work on a cluster computing project myself?

How to get more done …

CS514

- Work Harder
- Work Faster
- Get Help
- Processor Speed
- Algorithms
- Parallel processing

Some History

CS514

- Von Neuman thought parallelism to be impossible
- ILLIAC IV – first massive parallel machine (Illinois '60)
- Japan's 5th Generation Project
- USA – Grand Challenges
- Commercial: NCR, IBM Fijustu, Intel SSD, Gray, Convex

Traditional Users

CS514

- Scientists – investigate the unknown
- Engineers – simulations
- Retailers – data mining
- Airlines – how to overbook
- Financial – gaining 0.1% advantage
- ……

The collapse of the Supercomputing Industry

CS514

- ’97 the industry icon Cray Research went bankrupt.
- Many reasons were given, among which the end of the cold war

The Real Reasons

CS514

- Microprocessors got fast, a lot faster
- High Availability became a mass market.

Some early examples

CS514

- Vax-Clusters / OpenVMS Cluster

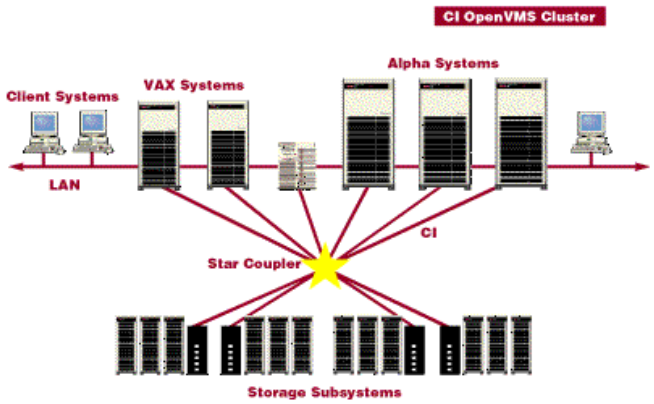

Tandem Himalaya

CS514

- Traditional for the HA – market

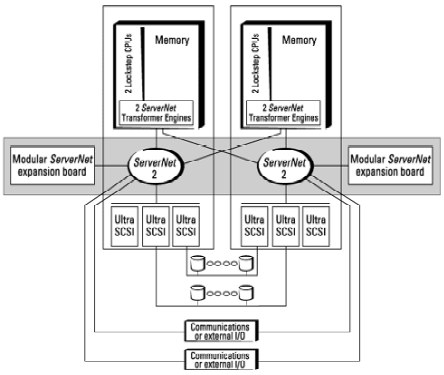

IBM SYSPLEX

CS514

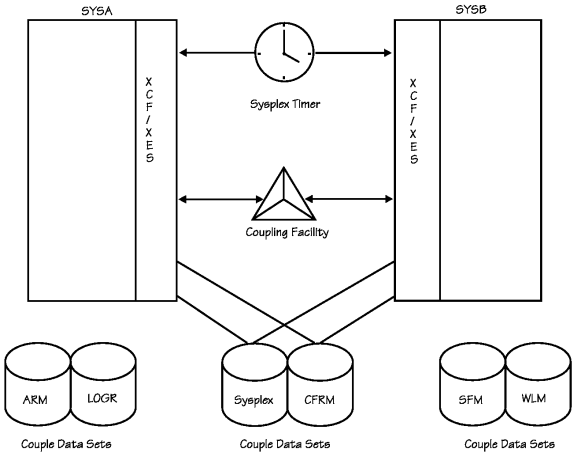

Cluster Definition

CS514

- Consists of a collection of interconnected whole computers
- Is used as a single, unified computing resource

Distinction from Other Systems

CS514

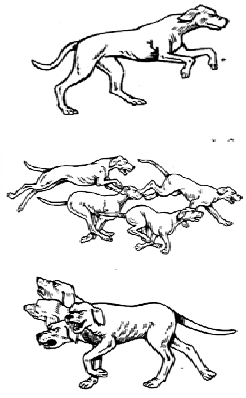

- **Scaling**: Adding a head or a whole dog
- **Availability**: what if a dog breaks a leg?
- **System management**: walking the dog
- **Software licensing**: dog tax

Technical Challenges - I

CS514

- Cluster Hardware (NOW, rack&stack, NUMAs)
- Cluster Communication (Interconnects, Communication Protocols)
- Cluster System Middleware (management, availability, tools)
- High-performance IO systems (storage, file systems, data placement and movement)

Technical Challenges - II

CS514

- Job and Resource Management
- Programming Environments (Distr. Objects, Message Passing)
- Scalable Services.
- Business frameworks (multi tiers, web based, decision support)
- Applications (Scientific, High-Availability, Scalable performance)

Single System Image

CS514

- From the perspective of
  - User
  - Network
  - Application
  - Administrator
- Key Issues:
  - Each SSI has a boundary
  - SSI support can exist at different levels

Single System Image - II

CS514

- Boundary:
  - Inside a single machine
  - Outside a collection of machines
- SSI Levels
  - Application
  - Middleware
  - Operating System
  - Hardware

Is Transparency a *good thing*?

CS514

- Yes, but achieving it is close to impossible
- Many transparencies were introduced in legacy code with disastrous side effects
- User to cluster is possible
- Server side should be avoided

1997 Microsoft.Com Web site (4 farms).

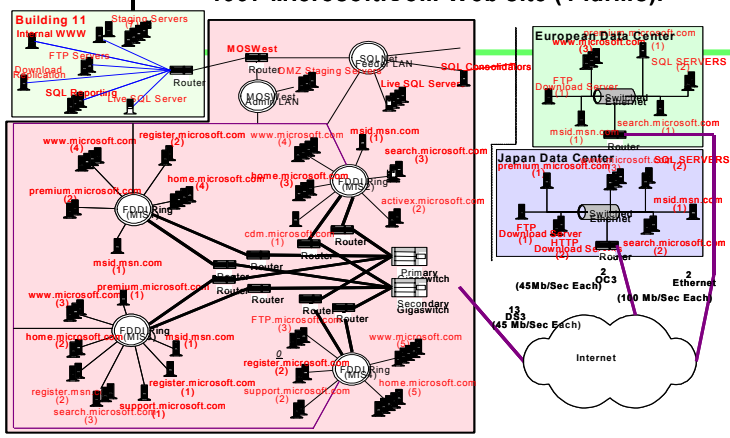

Two main Cluster Types

CS514

- RACS – *Reliable Array of Cloned Services*

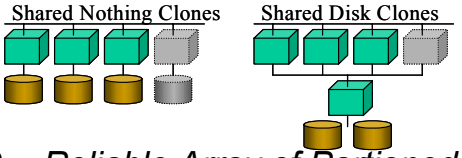

- RAPS – *Reliable Array of Partioned Services*

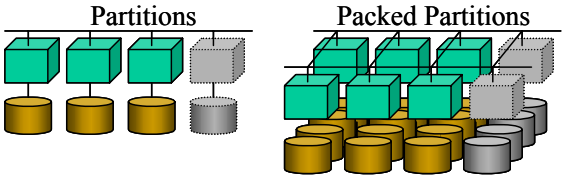

Taxonomy

CS514

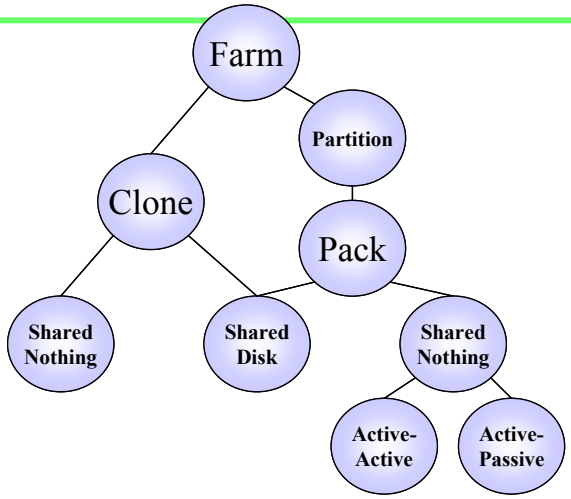

Sample E-Commerce Farm

The FARM: Clones and Packs of Partitions

CS514

Packed Partitions: Database Transparency

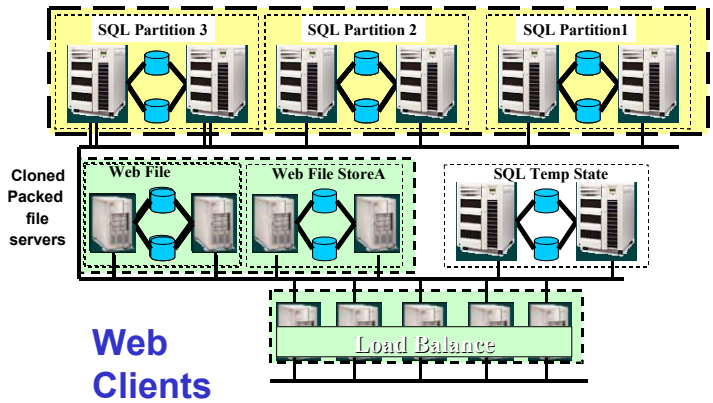

Windows NT Clusters - MSCS

CS514

- Group of independent systems that appear as a single system
- Managed as a single system
- Common namespace
- Services are “cluster-wide”
- Ability to tolerate component failures
- Components can be added transparently to users
- Existing client connectivity is not affected by clustered applications

MSCS Features

CS514

- Shared nothing
  - Simplified hardware configuration
- Remoteable tools
- Windows NT manageability enhancements
  - Never take a “cluster” down: shell game rolling upgrade
- Microsoft® BackOffice™ product support
- Provide clustering solutions for all levels of customer requirements
  - Eliminate cost and complexity barriers

Non-Features Of MSCS

CS514

- Not lock-step/fault-tolerant
- Not able to “move” running applications
  - “MSCS” restarts applications that are failed over to other cluster members
- Not able to recover shared state between client and server (i.e., file position)
  - All client/server transactions should be atomic
  - Standard client/server development rules still apply
    - Atomic Consistent Isolated Durable (ACID) always wins

Basic MSCS Terms

CS514

- Quorum Resource
  - Usually (but not necessarily) a SCSI disk
  - Requirements:
    - Arbitrates for a resource by supporting the challenge/defense protocol
    - Capable of storing cluster registry and logs
  - Used to Persist Configuration Change Logs
    - Tracks changes to configuration database when any defined member missing (not active)
    - Prevents configuration partitions in time
      - "Temporal Partitions"

"MSCS" Cluster

CS514

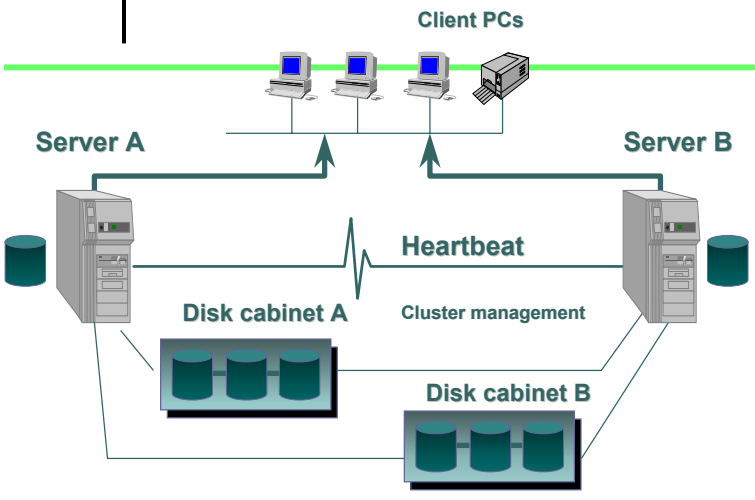

"MSCS" Namespace

Cluster view

CS514

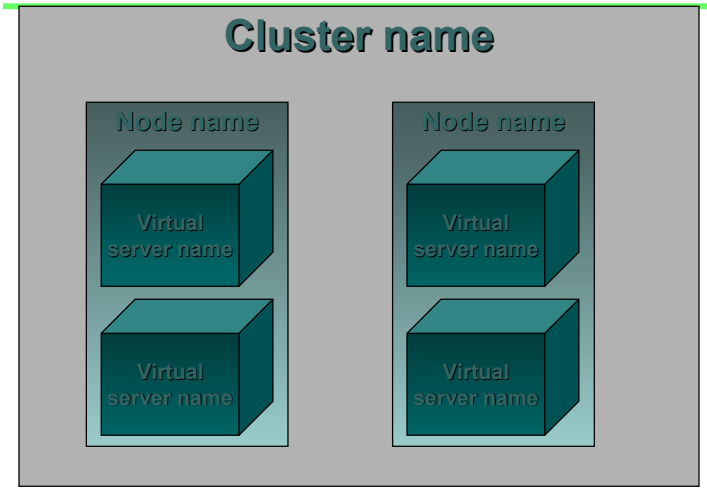

MSCS Namespace

Outside world view

CS514

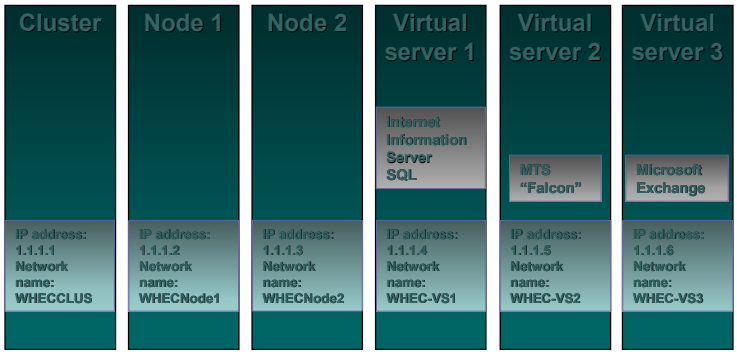

MSCS Architecture

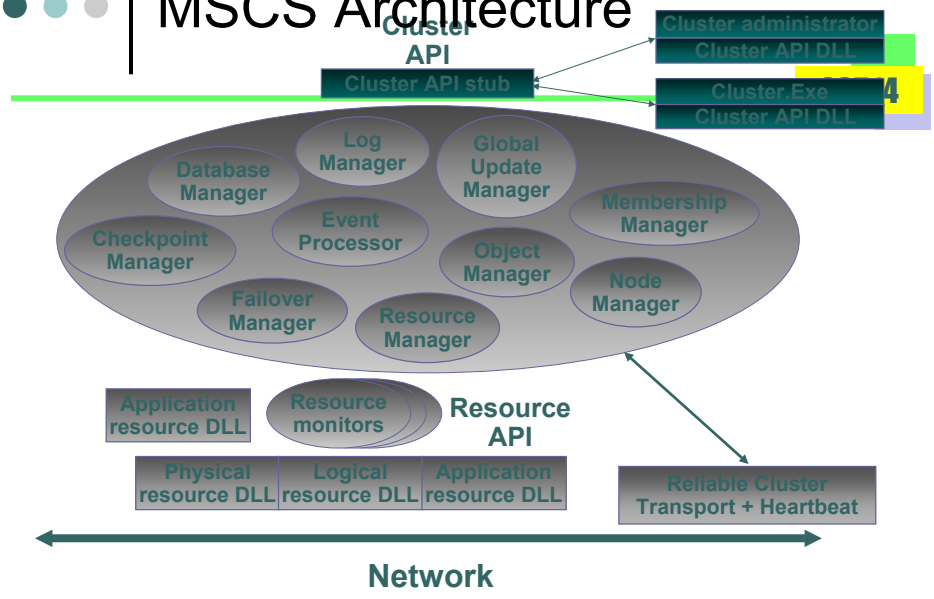

MSCS Architecture

CS514

- Cluster service is comprised of the following objects
  - Failover Manager (FM)
  - Resource Manager (RM)
  - Node Manager (NM)
  - Membership Manager (MM)
  - Event Processor (EVT)
  - Database Manager (DM)
  - Object Manager (OM)
  - Global Update Manager (GUM)
  - Checkpoint Manager (CM)

Cluster Communications

CS514

- Most communication via RPC
- UDP used for membership heartbeat messages
- Standard (e.g., Ethernet) interconnects

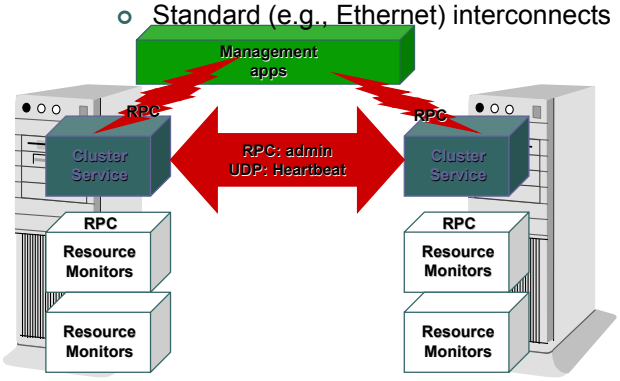

MSCS Architecture

- Top tier provides cluster abstractions
- Middle tier provides distributed operations
- Bottom tier is Windows NT and drivers

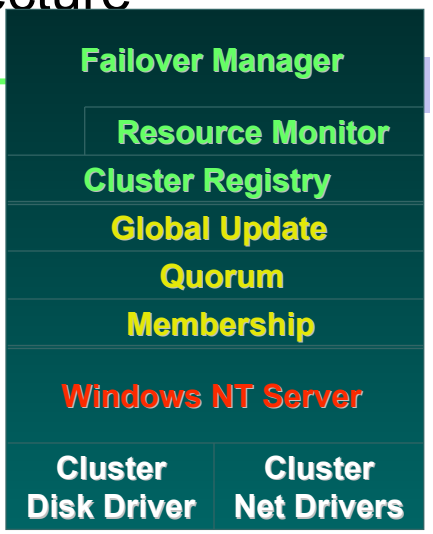

Membership And Regroup

CS514

- Membership:
  - Used for *orderly* addition and removal from { active nodes }
- Regroup:
  - Used for failure detection (via heartbeat messages)
  - *Forceful* eviction from { active nodes }

Membership

CS514

- Defined cluster = all nodes
- Active cluster:
  - Subset of defined cluster
  - Includes Quorum Resource
    - Transitive Ownership
  - Stable (no regroup in progress)

Challenge/Defense Protocol

CS514

- SCSI-2 has reserve/release verbs
  - Semaphore on disk controller
- Owner gets lease on semaphore
- Renews lease once every 3 seconds
- To preempt ownership:
  - Challenger clears semaphore (SCSI bus reset)
  - Waits 10 seconds
    - 3 seconds for renewal + 2 seconds bus settle time
    - x2 to give owner two chances to renew
  - If still clear, then former owner loses lease
  - Challenger issues reserve to acquire semaphore

Challenge/Defense Protocol:

Successful Defense

CS514

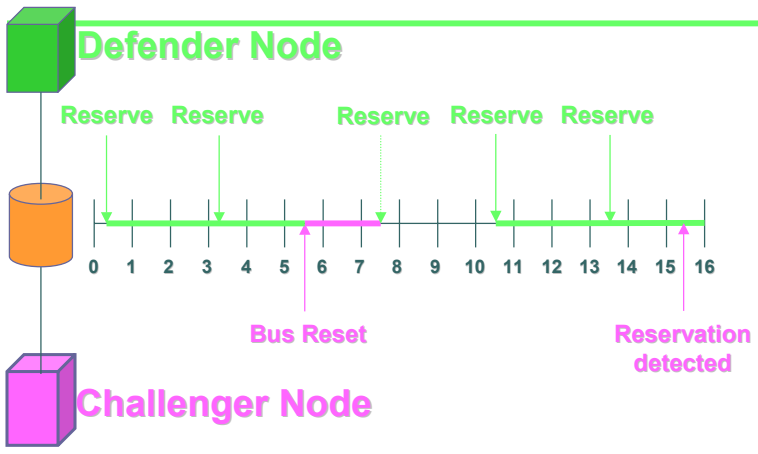

Challenge/Defense Protocol:

Successful Challenge

CS514

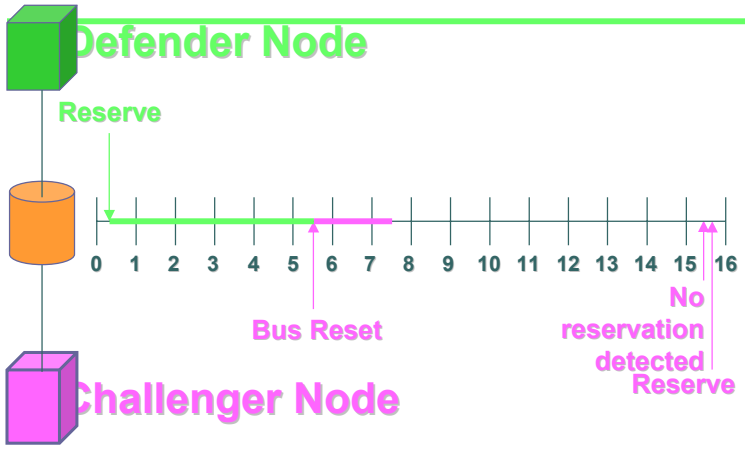

Regroup

CS514

- Invariant:
  All members agree on {members}
- *Regroup* re-computes {members}
- Each node sends *heartbeat* message to a peer (default is one per second)
- Regroup if two lost heartbeat messages
  - *Suspicion* that sender is dead
  - Failure detection in bounded time
- Uses a 5-round protocol to agree
  - Checks communication among nodes
  - Suspected missing node may survive
- Upper levels (global update, etc.) informed of regroup event

Membership State Machine

CS514

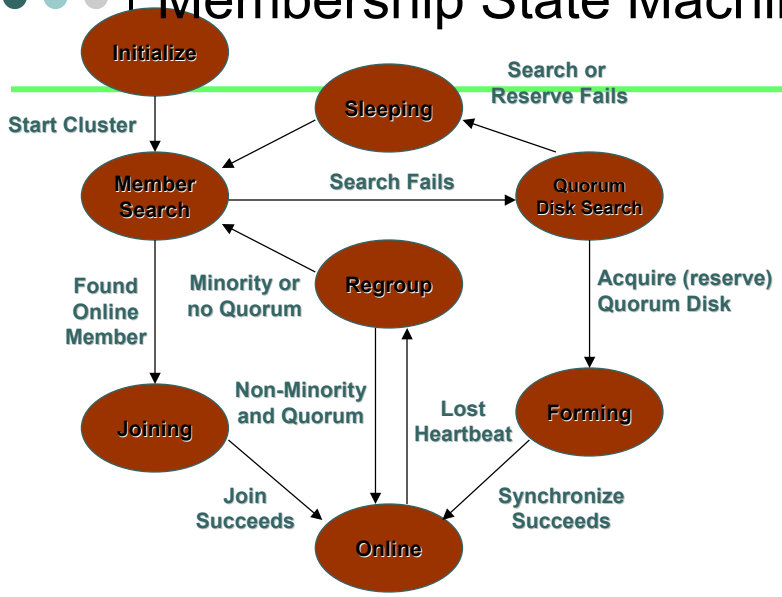

Joining A Cluster

CS514

- When a node starts up, it mounts and configures only local, non-cluster devices
- Starts Cluster Service which
  - Looks in local (stale) registry for members
  - Asks each member in turn to sponsor new node's membership. (Stop when sponsor found.)
- Sponsor (any active member)
  - Sponsor authenticates applicant
  - Broadcasts applicant to cluster members
  - Sponsor sends updated registry to applicant
  - Applicant becomes a cluster member

Forming A Cluster (When Joining Fails)

CS514

- Use registry to find quorum resource
- Attach to (arbitrate for) quorum resource
- Update cluster registry from quorum resource
  - E.g., if we were down when it was in use
- Form new one-node cluster
- Bring other cluster resources online
- Let others join your cluster

Leaving A Cluster (Gracefully)

CS514

- Pause:
  - Move all groups off this member.
  - Change to paused state (remains a cluster member)
- Offline:
  - Move all groups off this member.
  - Sends ClusterExit message all cluster members
    - Prevents regroup
    - Prevents stalls during departure transitions
  - Close Cluster connections (now not an active cluster member)
  - Cluster service stops on node
- Evict: remove node from defined member list

Leaving A Cluster (Node Failure)

CS514

- Node (or communication) failure triggers Regroup
- If after regroup:
  - Minority group OR no quorum device:
    - Group does NOT survive
  - Non-minority group AND quorum device:
    - Group DOES survive
- Non-Minority rule:
  - Number of new members >= 1/2 old *active* cluster
  - Prevents minority from seizing quorum device at the expense of a larger potentially surviving cluster
- Quorum guarantees *correctness*
  - Prevents “split-brain”
    - *E.g., with newly forming cluster containing a single node*

Global Update

CS514

- Propagates updates to all nodes in cluster
- Used to maintain replicated cluster registry
- Updates are atomic and totally ordered
- Tolerates all benign failures.
- Depends on membership
  - All are up
  - All can communicate
- R. Carr, Tandem Systems Review. V1.2 1985, sketches regroup and global update protocol

Global Update Algorithm

CS514

- Cluster has locker node that regulates updates
  - Oldest active node in cluster
- Send Update to locker node
- Update other (active) nodes
  - In seniority order (e.g. locker first)
  - This includes the updating node
- Failure of all updated nodes:
  - Update never happened
  - Updated nodes will roll back on recovery
- Survival of any updated nodes:
  - New locker is oldest and so has update if any do
  - New locker restarts update

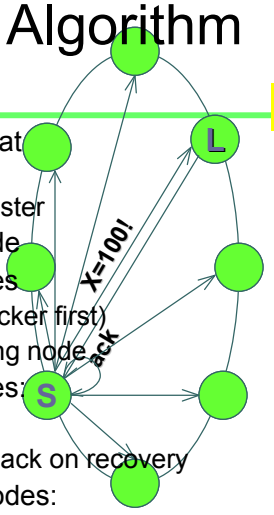

Cluster Registry

CS514

- Separate from local Windows NT Registry
- Maintains cluster configuration
  - Members, resources, restart parameters, etc.
- Stable storage
- Replicated at each member
  - Global Update protocol
  - Windows NT Registry keeps local copy

Cluster Registry Bootstrapping

CS514

- Membership uses Cluster Registry for list of nodes
  - *…Circular dependency*
- Solution:
  - Membership uses stale local cluster registry
  - Refresh after joining or forming cluster
  - Master is either
    - Quorum device, or
    - Active members

Resource Monitor

CS514

- Polls resources:
  - IsAlive and LooksAlive
- Detects failures
  - polling failure
  - failure event from resource
- Higher levels tell it
  - Online, Offline
  - Restart

Failover Manager

- Assigns groups to nodes based on
  - Failover parameters
  - Possible nodes for each resource in group
  - Preferred nodes for resource group

Resource

CS514

- Fails over (moves) from one machine to another
  - Logical disk
  - IP address
  - Server application
  - Database
- Online at only one machine
- May depend on another resource
- Well-defined properties controlling its behavior

Resource Properties

CS514

- Resource type
- Poll intervals
  - Looksalive
  - Isalive
- Private resource data
  - Unique identifier
  - Hardware binding
- Group membership
- Possible nodes
- Restart policy
- Dependencies

Time

CS514

- Time must increase monotonically
  - Otherwise applications get confused
  - E.g., make/nmake/build
- Time is maintained within failover resolution
  - Not hard, since failover on order of seconds
- Time is a resource, so one node owns time resource
- Other nodes periodically correct drift from owner's time

Win 2000 MSCS Cluster

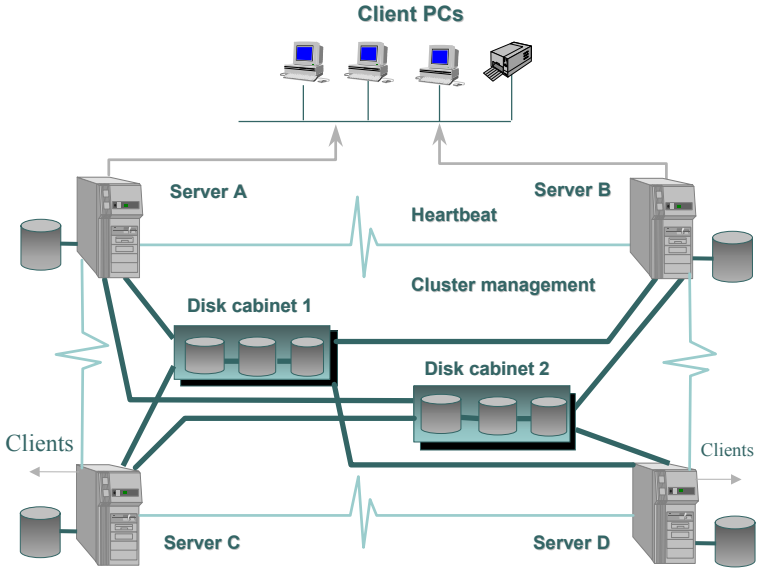

Application Evolution

MSCS Version NT 4.0

CS514

Application	Node 1	Node 2
Microsoft SQL Server	☑	
Microsoft Transaction Server (MTS)		☑
Internet Information Server (IIS)	☑	
Microsoft Exchange Server		☑

Application Evolution

MSCS Version Windows 2000

CS514

Application	Node 1	Node 2	Node 3	Node 4
Microsoft SQL Server	☑	☑	☑	☑
Microsoft Transaction Server (MTS)	☑	☑	☑	☑
Internet Information Server (IIS)	☑	☑	☑	☑
Microsoft Exchange Server	☑	☑	☑	☑