Understanding VMScluster State Transitions

Keith Parris
VMS Engineering

Why have state transitions?

Figure 1: Foundation for shared access

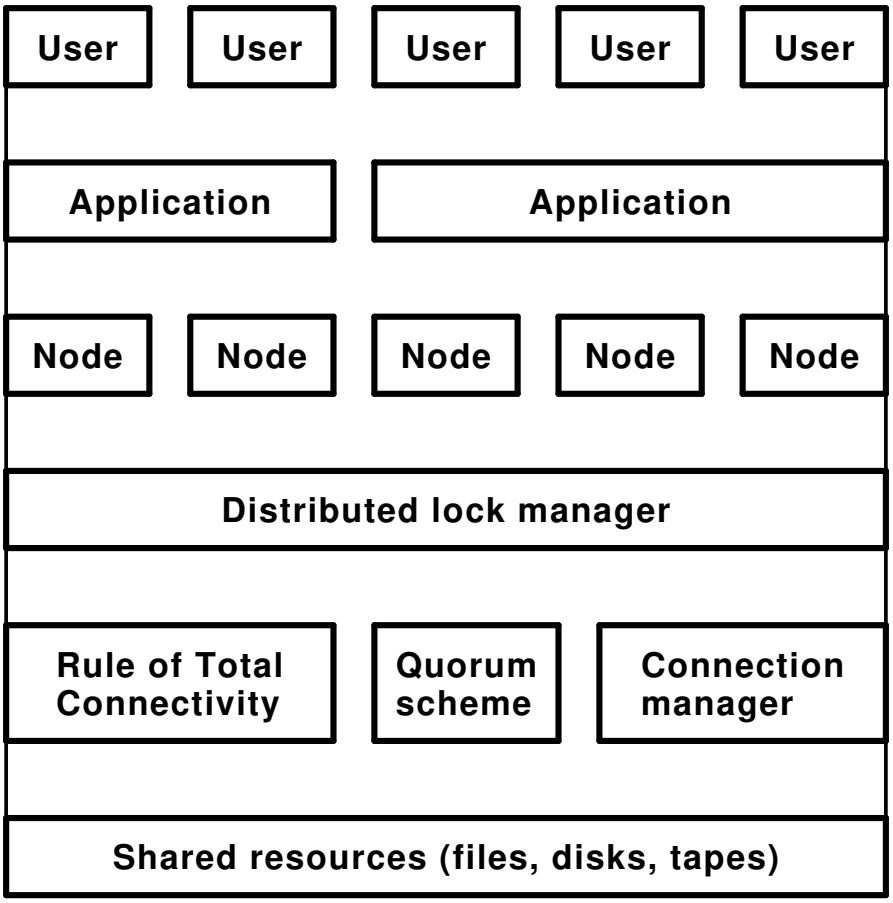

ART_FOUNDATION

Distributed lock manager

Figure 2: Distributed locks

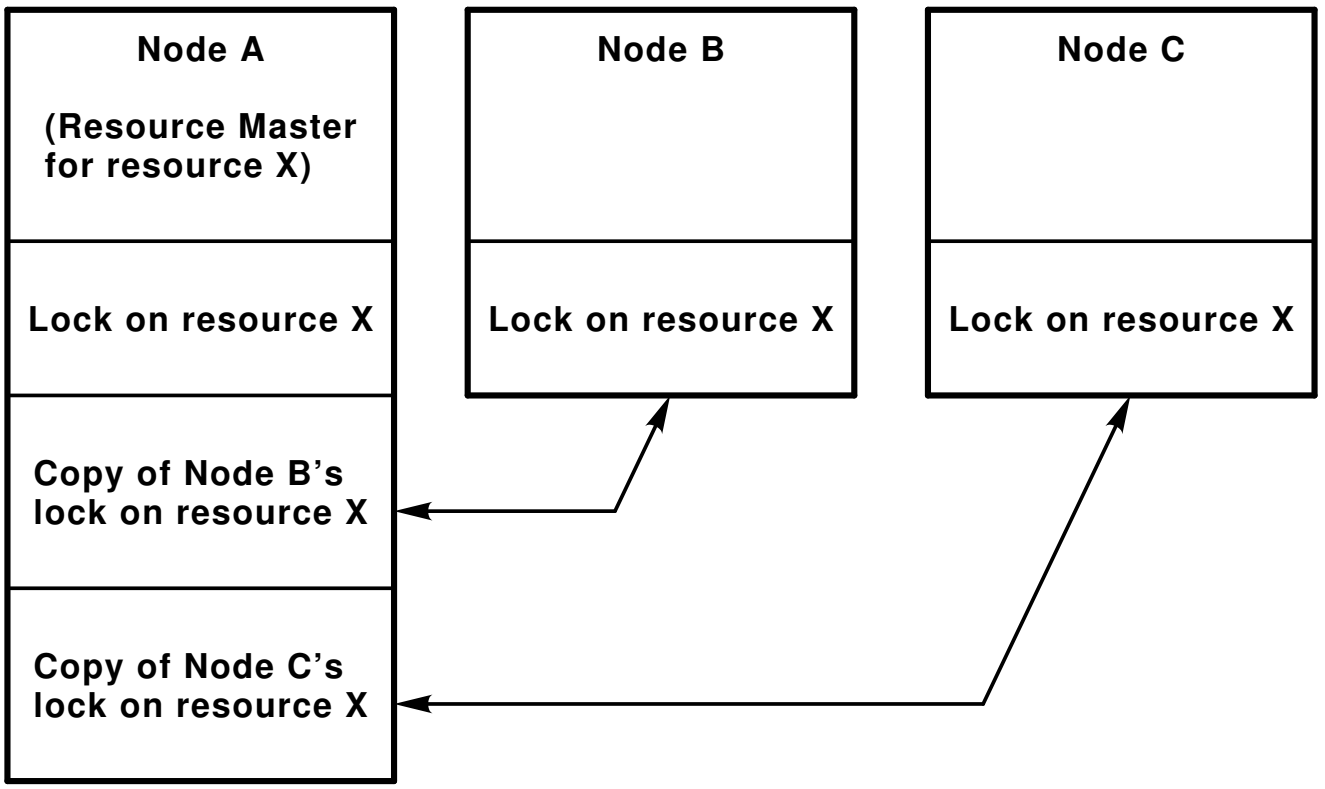

ART_DISTR_LOCKS

Distributed lock manager

Figure 3: Directory lookup

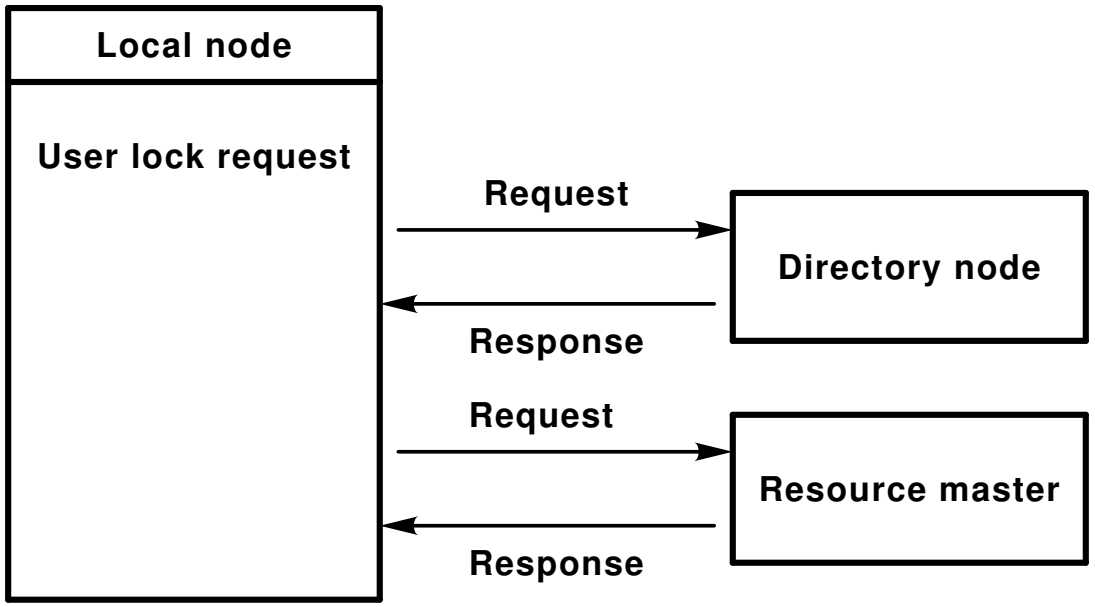

Possible responses to directory lookup:

1. Request granted (directory node is also resource master)
2. Become master node (no other locks on this resource)
3. Node *xxx* is the resource master

Distributed lock manager

Figure 4: Lock directory weights and resulting vector

Big node:	PAPABR	Lock Directory Weight = 2
Middle–size node:	MAMABR	Lock Directory Weight = 1
Satellite node:	BABYBR	Lock Directory Weight = 0

Resulting lock directory vector:

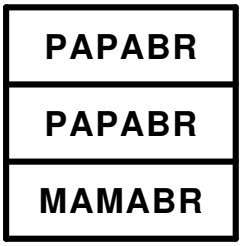

Implications for state transitions:

Node leaving or joining	Effect on directory vector
LOCKDIRWT zero:	No effect
LOCKDIRWT non–zero:	Must change

ART_LOCKDIRWT

Distributed lock manager

Lock rebuild types

1. **Merge:**
   Enter already-held lock tree into directory on directory node. (Lock requests not stalled)
2. **Partial:**
   Release locks held by departed node, and remaster any locks it was resource master for.
3. **Directory:**
   Do what a Partial rebuild does, and also rebuild the directory information.
4. **Full:**
   Throw away all master copies of locks, select new resource masters, and then re-acquire locks, rebuilding the directory information in the process.

Node's LOCKDIRWT	State change	Type of rebuild
zero	Node joining	Merge
zero	Node leaving	Partial
non-zero	Node joining or leaving	Directory
n/a	Cluster formation	Full

Recent changes to lock manager state transition behavior:

- V5.0
  - Partial and merge rebuilds
- V5.2
  - Dynamic remastering
- V5.4
  - Remastering during shutdown
  - Multi-threaded lock rebuild
- V5.5
  - Load-based remastering

State transition triggers

A state transition is needed when:

- **A new member is added to the cluster**
- **An existing member is removed from the cluster**

New member detection: CI/DSSI

```
PANUMPOLL = 2
PAMAXPORT = 3
PAPOLLINTERVAL = 5 seconds
```

Figure 5: Polling on CI and DSSI

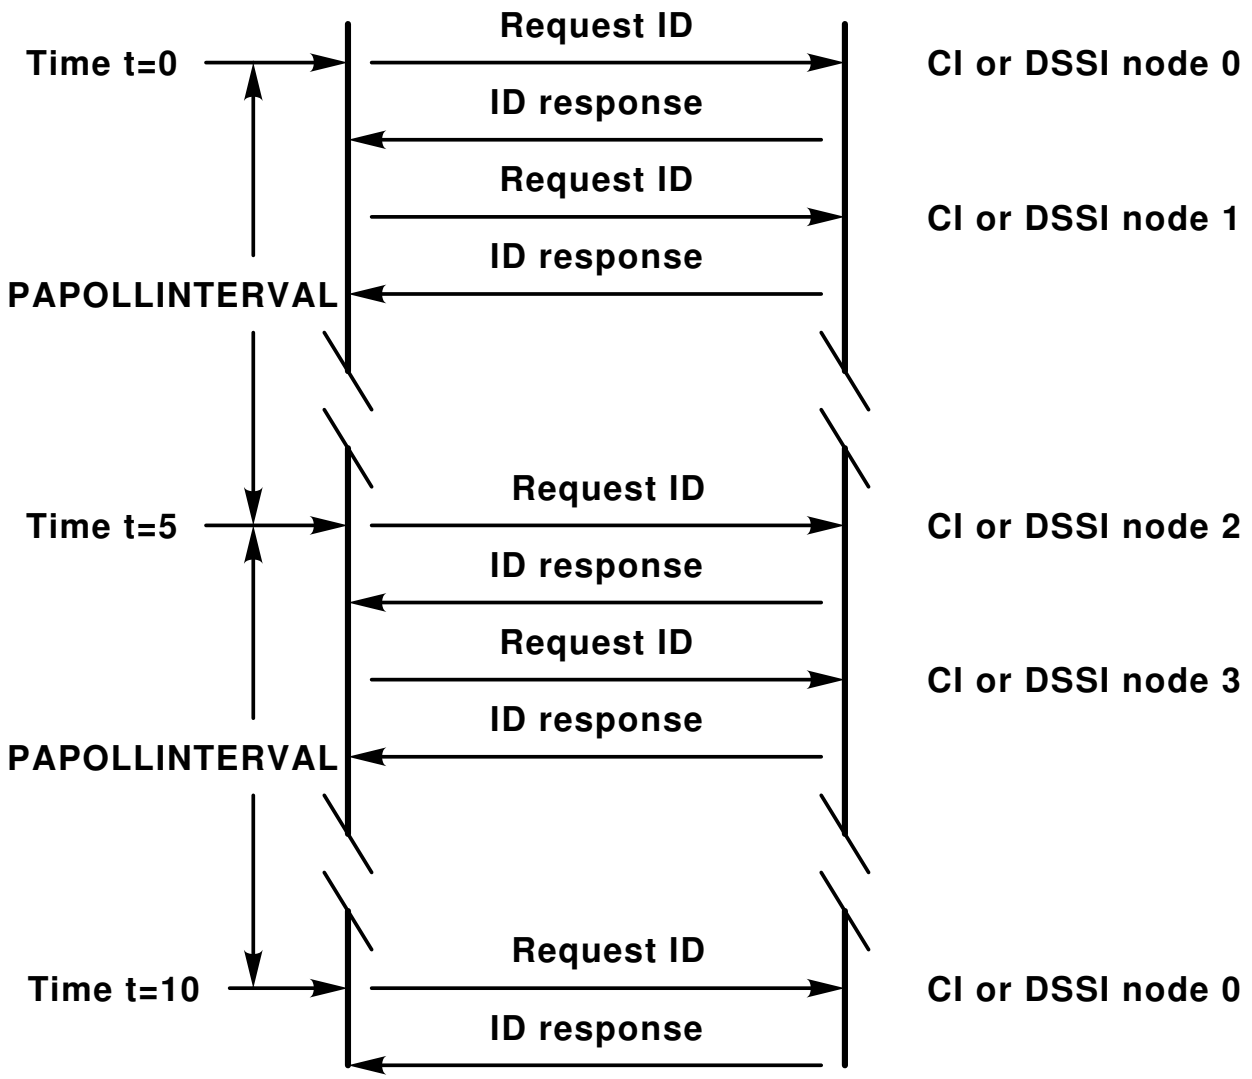

ART_CI_DSSI_POLL

New member detection: CI/DSSI

Figure 6: Virtual circuit formation on CI and DSSI

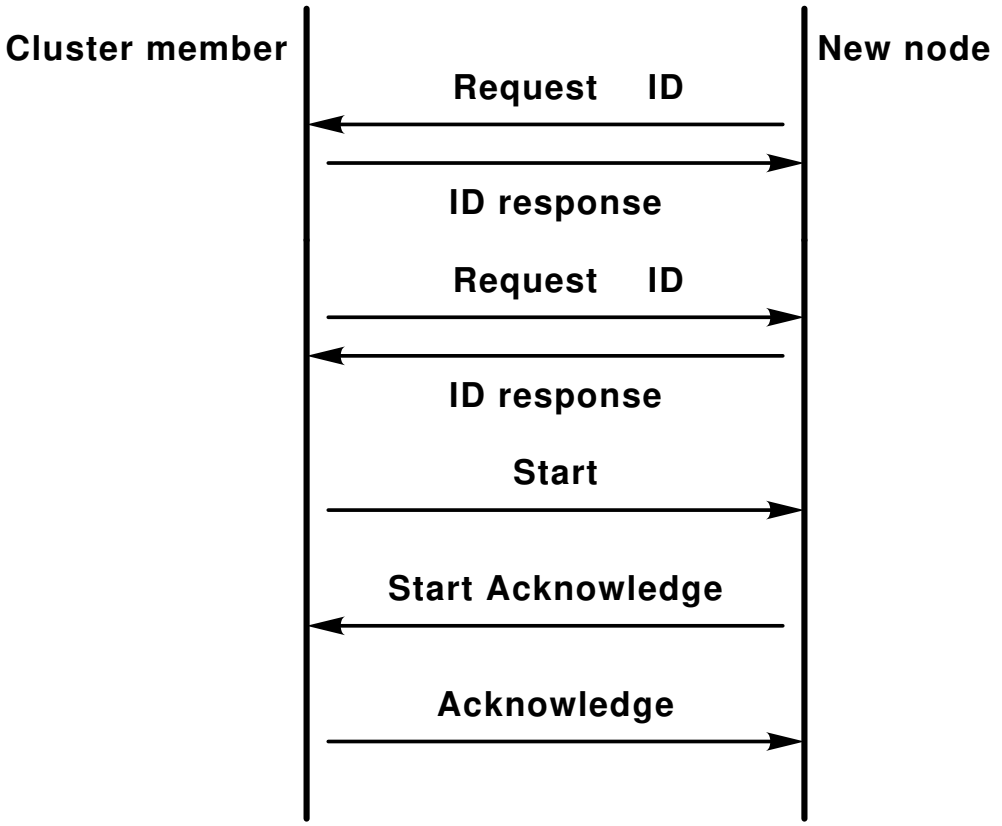

New member detection: Ethernet/FDDI

Figure 7: Virtual circuit formation on Ethernet and FDDI

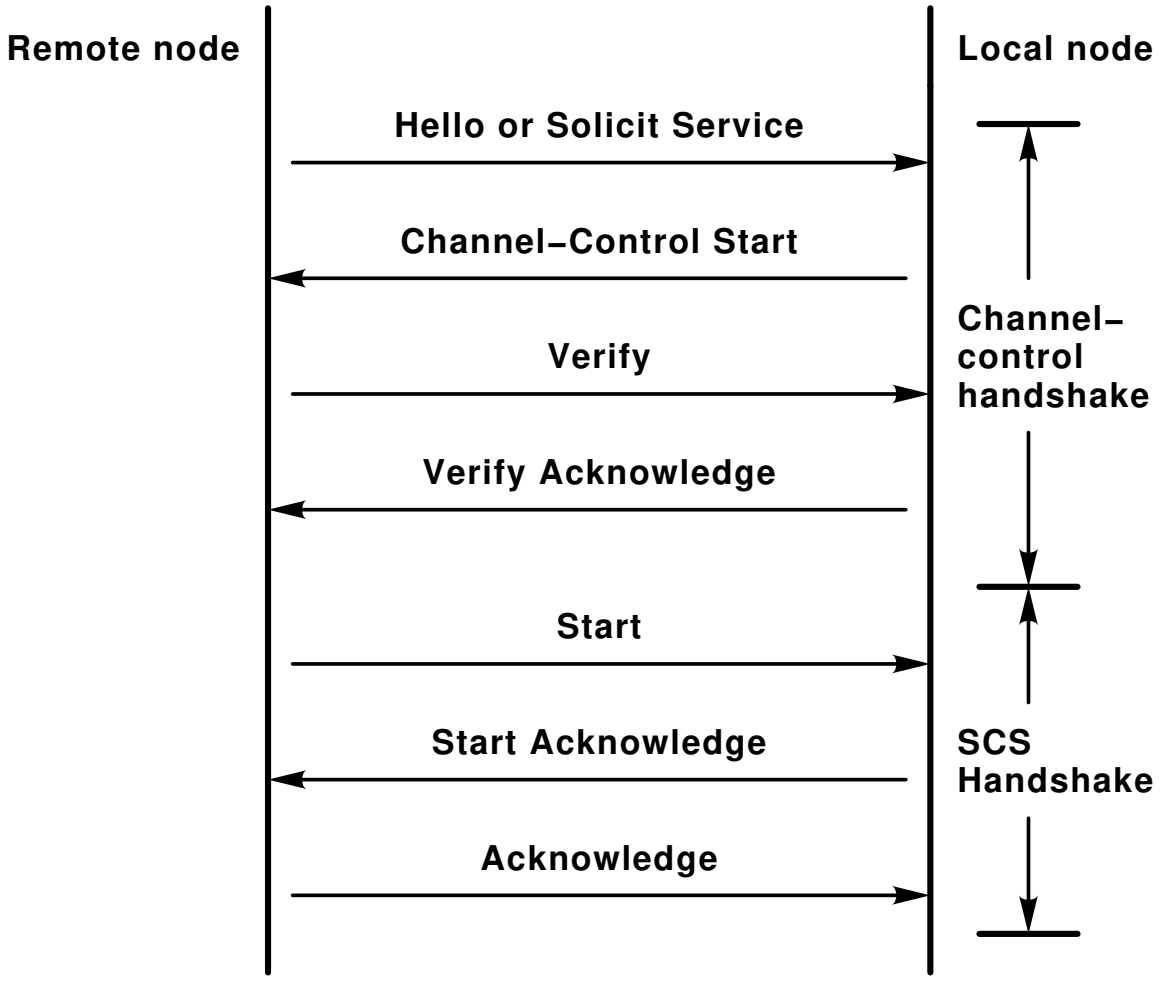

Failures

Types of failures:

- Node failure
- Adapter failure
- Interconnect failure

Failure examples:

- Thinwire station insertion (loss of a segment temporarily)
- Bridge loss and spanning-tree reconfiguration time
- Power failure for network equipment

Repair period

Recovery within repair period

Figure 8: Recovery within RECNXINTERVAL period

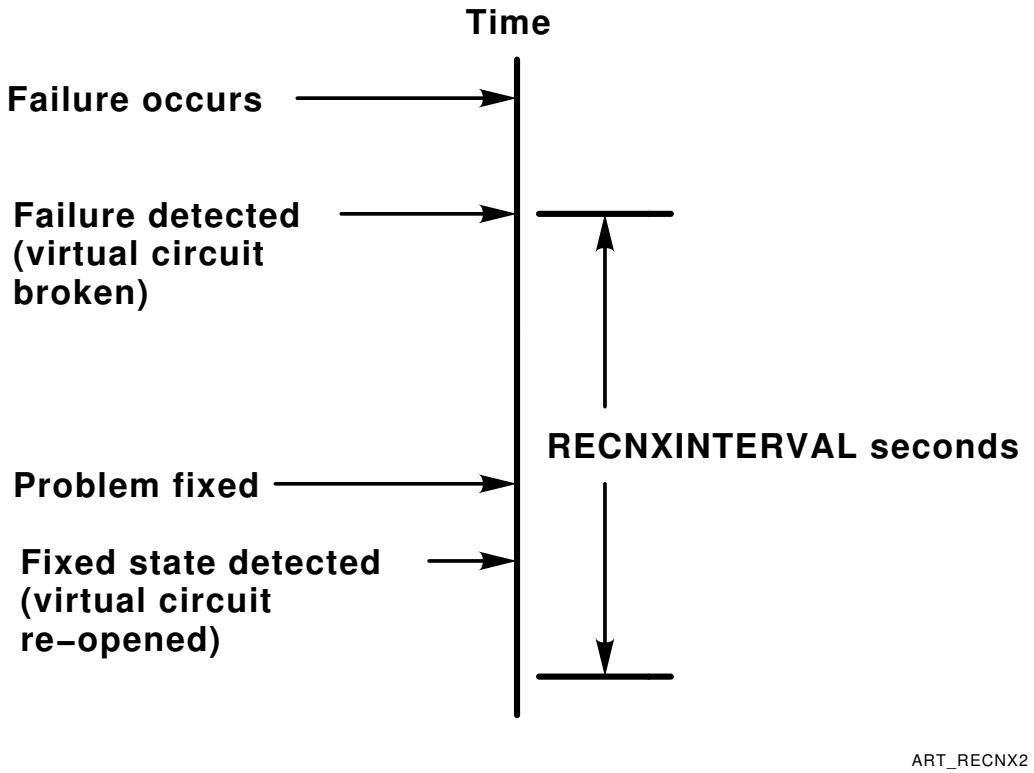

Repair period

Hard failure

Figure 9: Recovery period

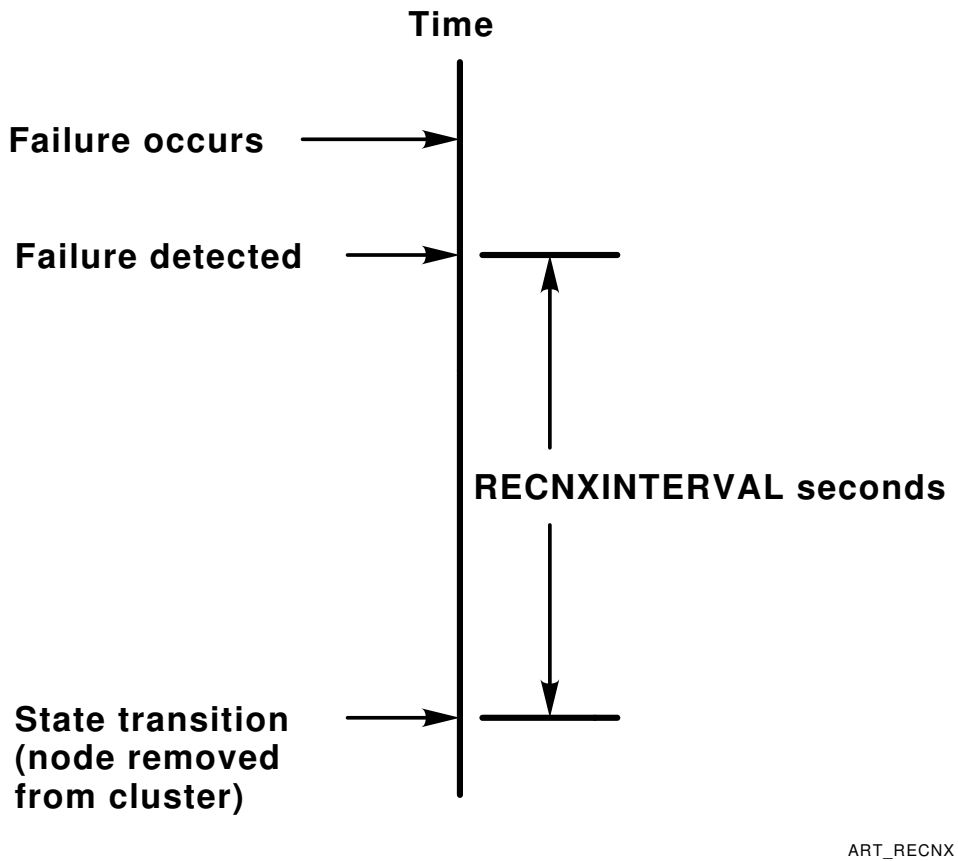

Repair period

Late recovery

Figure 10: Recovery after RECNXINTERVAL

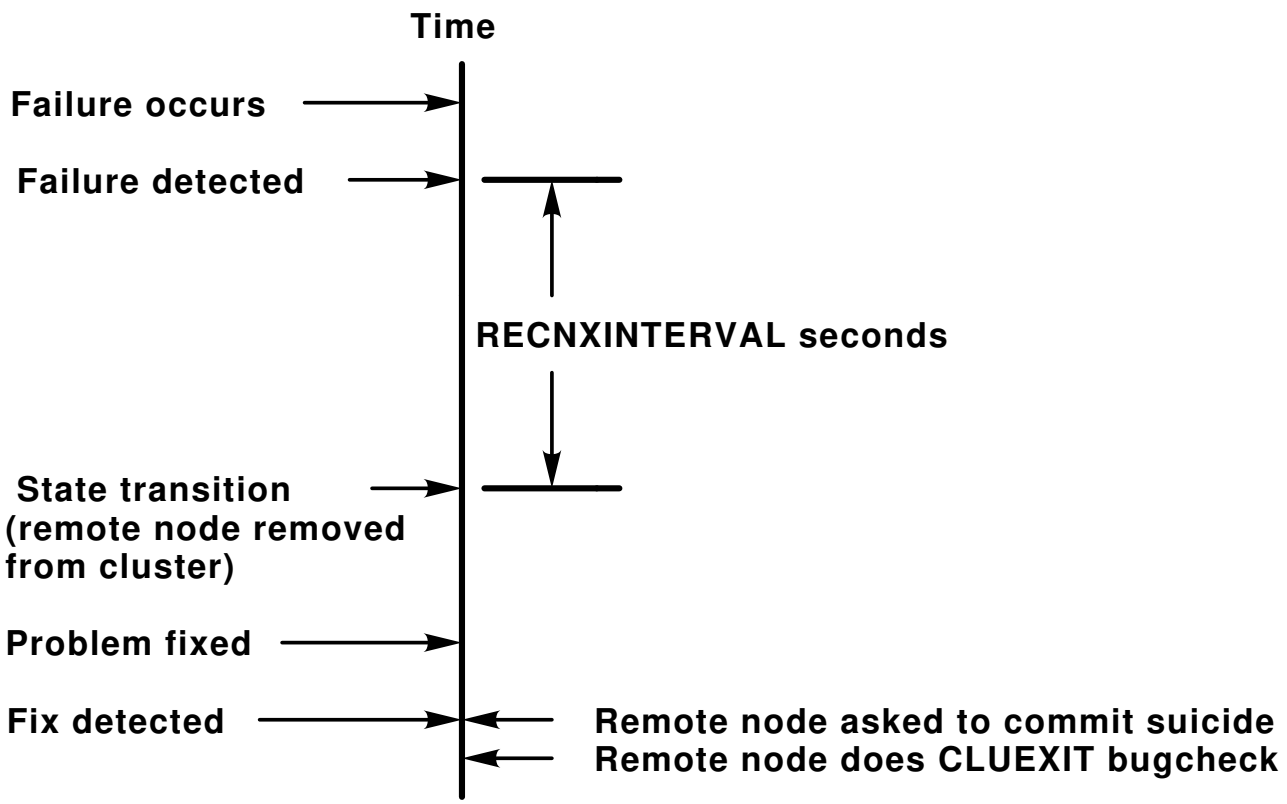

Failure detection mechanisms

Mechanisms to detect a node or communications failure:

- **Last-gasp datagram**
- **Periodic checking**
  - **— TIMVCFAIL check**
  - **— Multicast Hello messages for Ethernet and FDDI**
  - **— Polling for CI and DSSI**

Failure detection mechanisms

TIMVCFAIL check

Figure 11: Virtual-circuit check using TIMVCFAIL

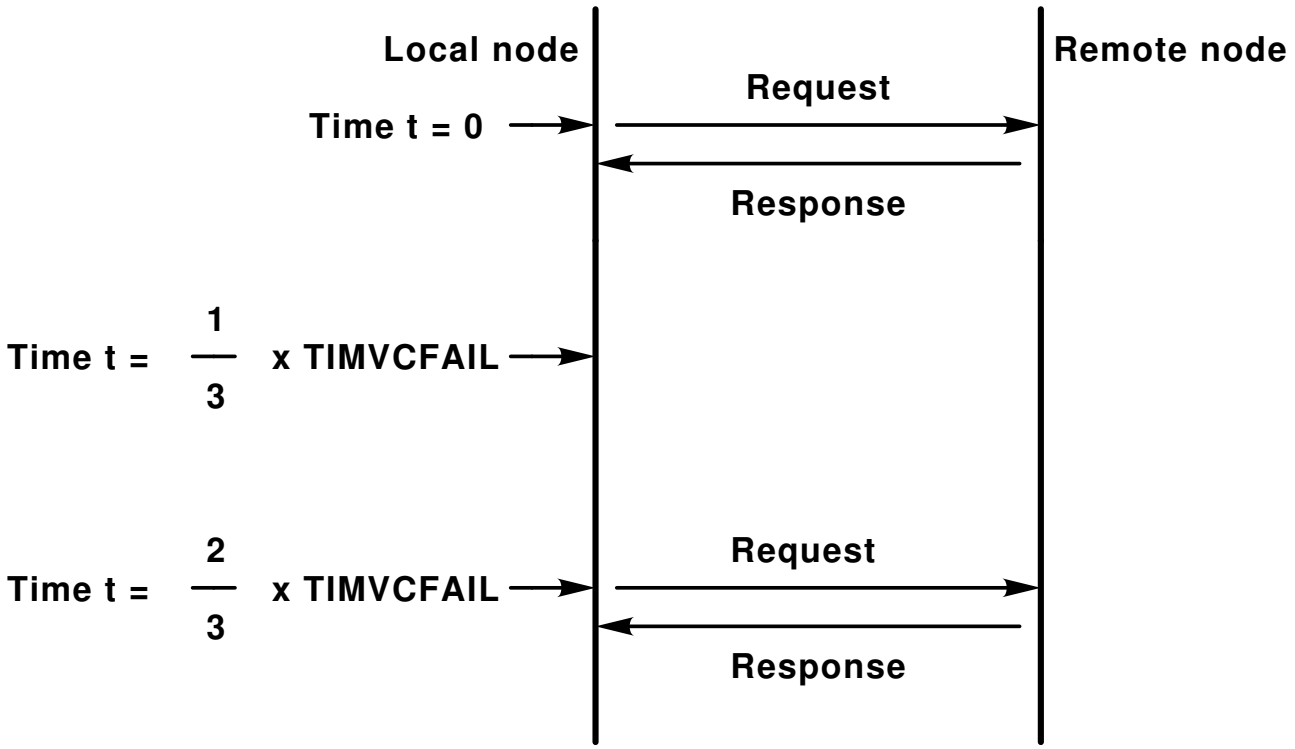

ART_TIMVC

Failure detection mechanisms

TIMVCFAIL failure detection

Figure 12: Failure detection using TIMVCFAIL

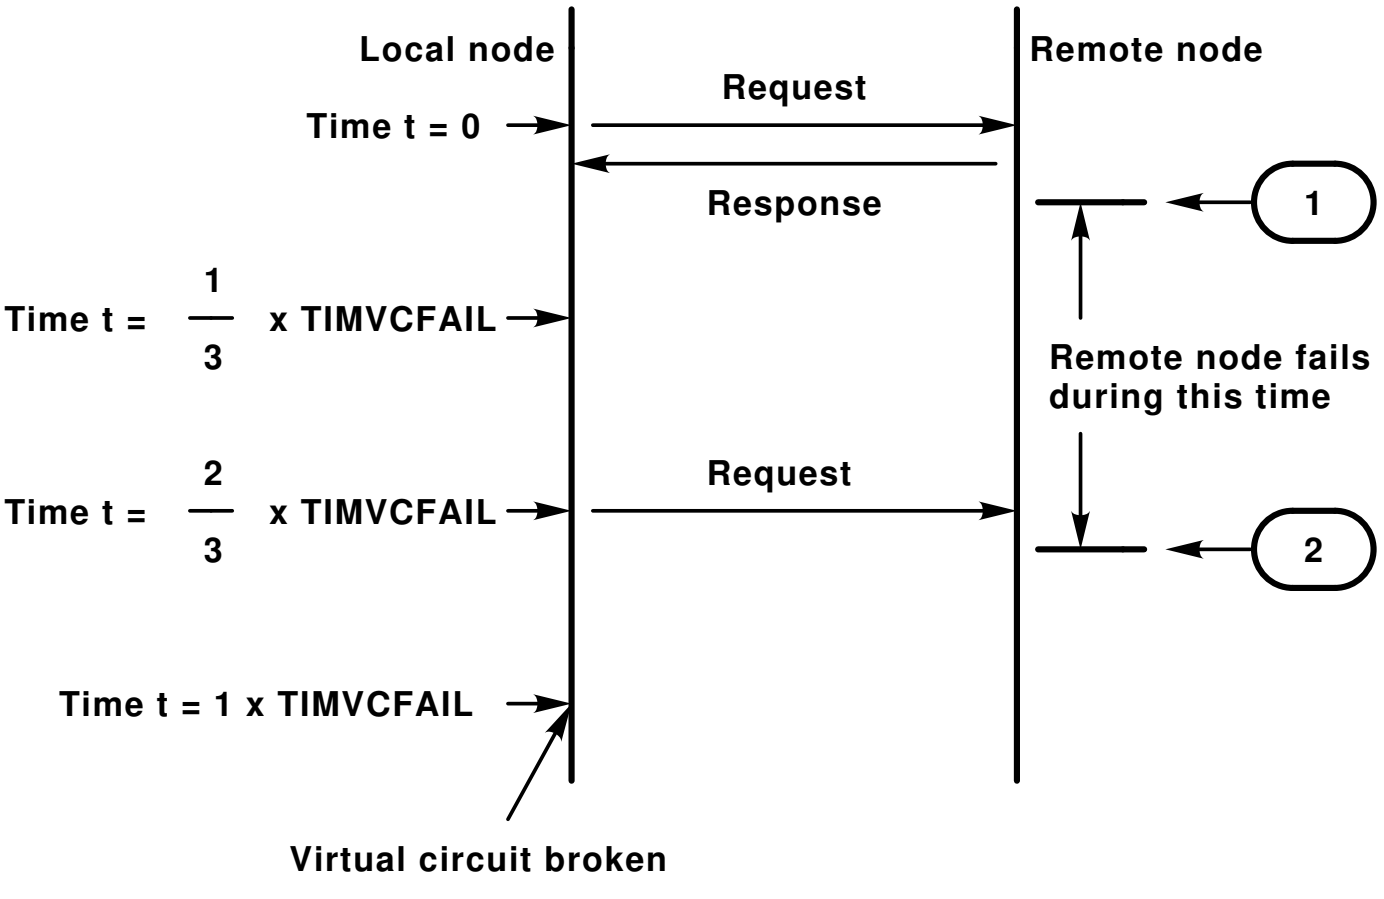

ART_TIMVC2

Failure detection mechanisms

Figure 13: Hello packets on LAN

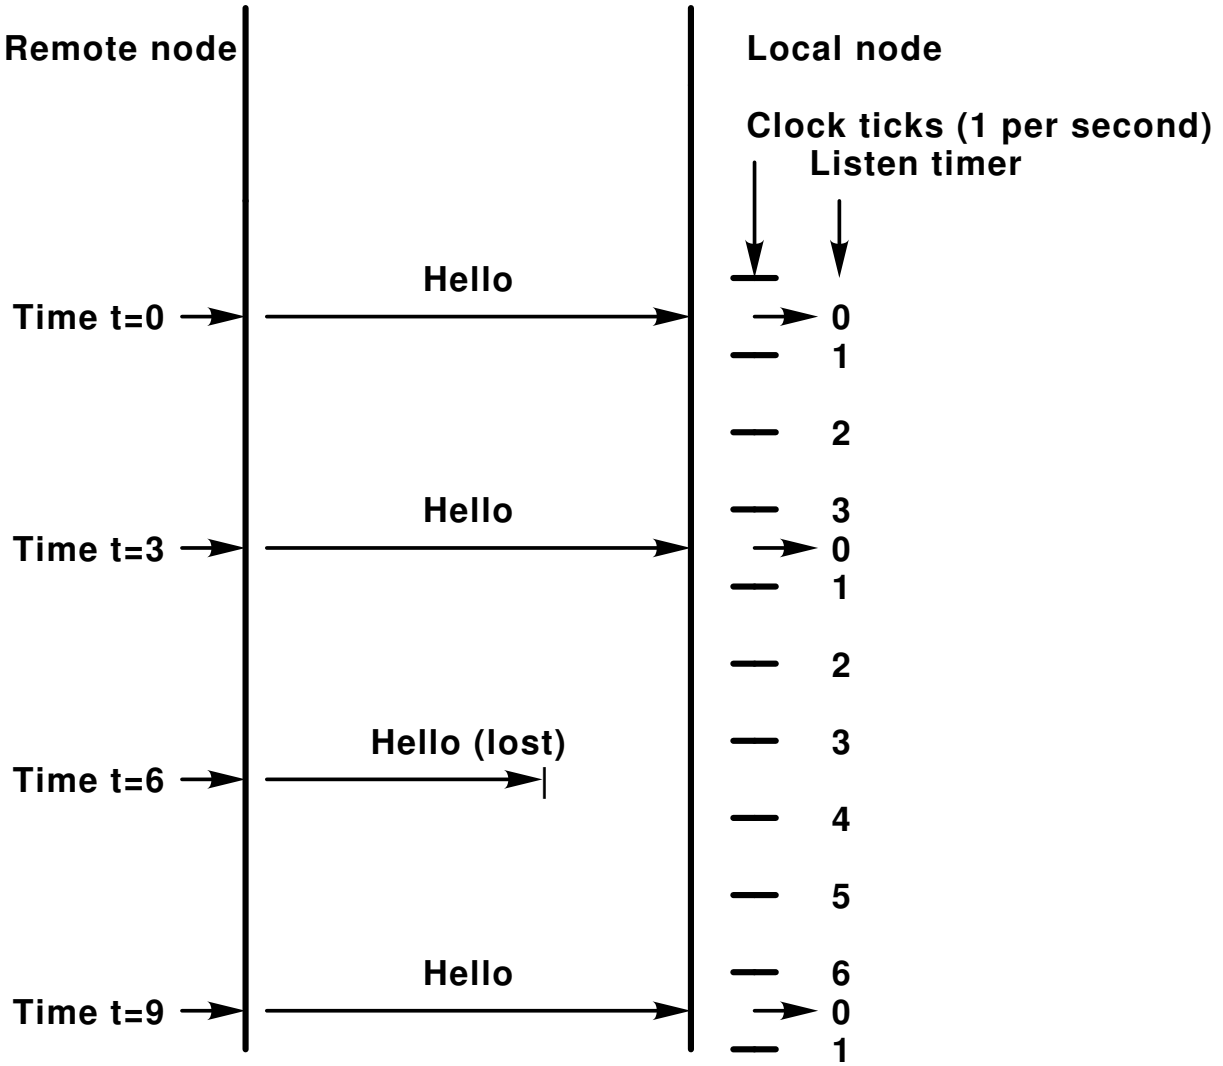

ART_HELLO

Failure detection mechanisms

Figure 14: Hello packets with failure

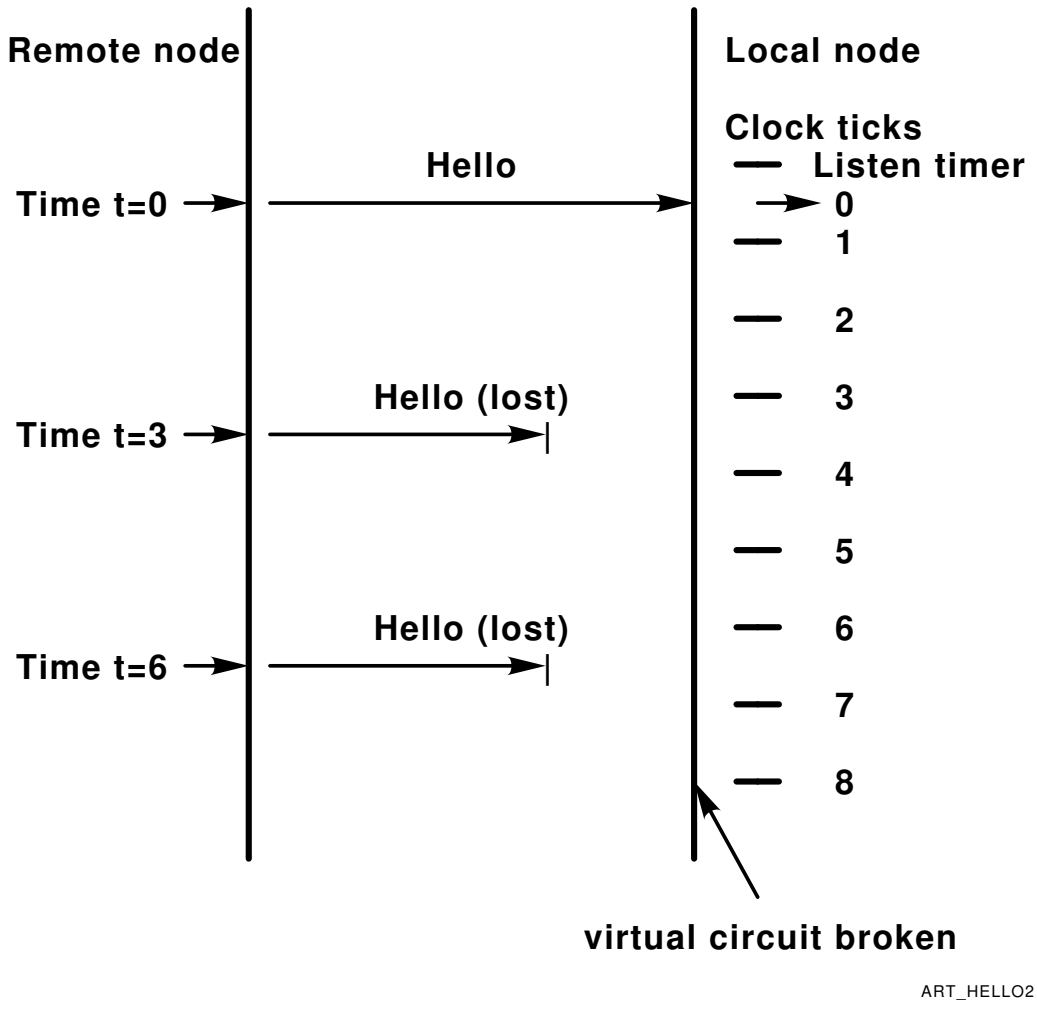

Avoiding state transitions

Single point of failure vs. redundancy

- **Multi-adapter CI/DSSI or Ethernet/FDDI**
- **Fail-over between CI/DSSI and Ethernet/FDDI**

Protection:

- **Power**
- **Network**

Effects of a state transition at user level

Quorum loss during quorum disk revalidation

Figure 15: Quorum disk revalidation

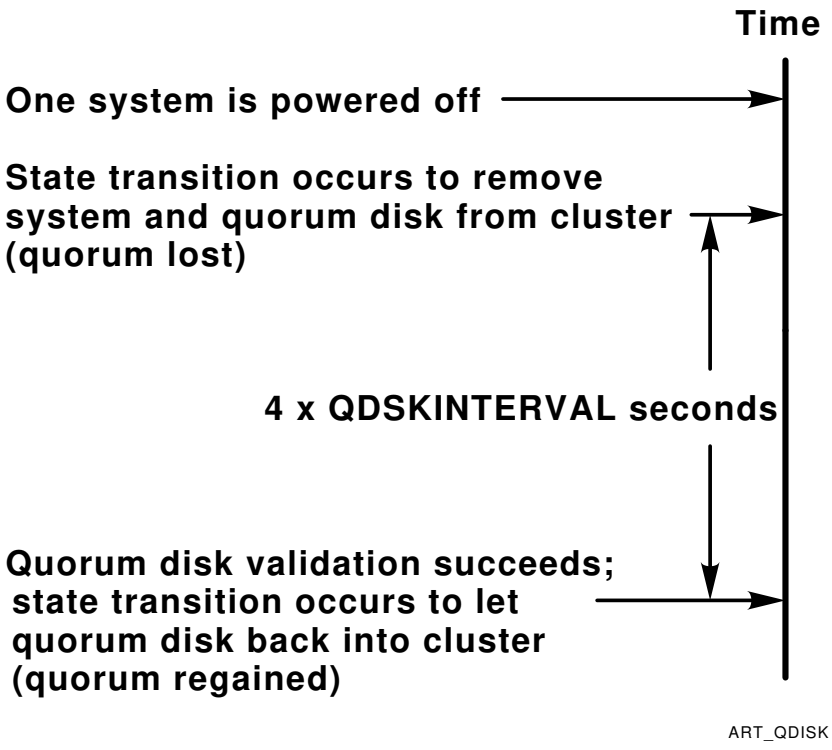

Effects of a state transition

Sequence of events

State transition tasks

1. Determine new cluster configuration.
   Quorum lost:
   - QUORUM capability bit removed from all CPUs
   - Disks put into mount verification

   Quorum not lost:
2. Rebuild lock database
   - Stall lock requests
   - I/O synchronization
   - Do rebuild work
   - Resume lock handling

Related issues

- Quorum disk validation
- Disk mount verification on quorum loss
- Node reboot activity
  - Startup work on rebooting big nodes
  - Satellites rebooting
  - Disk rebuilds
    * Set ACP_REBLDSYSD = 0
    * Use MOUNT/NOREBUILD
    * Do SET VOLUME/REBUILD in a batch job late at night
- Shadowing
- Application recovery work
  - Users dumped off a failed node log back in to another node
  - Replaying journal files

Measuring state transition effects

Determine the type of the last lock rebuild:

```
$ analyze/system
SDA> read sys$system:scsdef
SDA> evaluate @(@clu$gl_club + club$b_newrbld_req) & FF
Hex = 00000002   Decimal = 2  ACP$V_SWAPPRV
```

Rebuild types:

1. **Merge (locking not disabled)**
2. **Partial**
3. **Directory**
4. **Full**

Determine the duration of the last lock request stall period:

```
SDA> define toff = @(@clu$gl_club + club$l_toff)
SDA> define ton = @(@clu$gl_club + club$l_ton)
SDA> evaluate ton-toff
Hex = 0000026B   Decimal = 619  PDT$Q_COMQH+00003
```

Minimizing duration of state transitions

Configuration issues:

- **Few (e.g. 3) nodes**
- **Quorum VAX; no quorum disk**
- **Isolate cluster interconnects from other traffic**

Operational issues:

- **Avoid disk rebuilds on reboot**
- **Let Last-Gasp work**

Minimizing duration of state transitions

System parameters

- V5.4 or earlier:
  - PAPOLLINTERVAL = 1
  - PAMAXPORT = 31
  - PANUMPOLL = 32
  - PANOPOLL = 0 (leave polling enabled)
  - PASANITY = 1 (leave sanity timer enabled)
  - PASTIMOUT = 1
  - RECNXINTERVAL = 1
  - NISCS_LOAD_PEA0 = 0
- V5.5 or later:
  - TIMVCFAIL as low as 100 (1 second), as needed
  - RECNXINTERVAL as low as 1 second
  - NISCS_LOAD_PEA0 can be 0 or 1

Questions?