

Reality

Heartbeat FailSafe Feature

Introduction

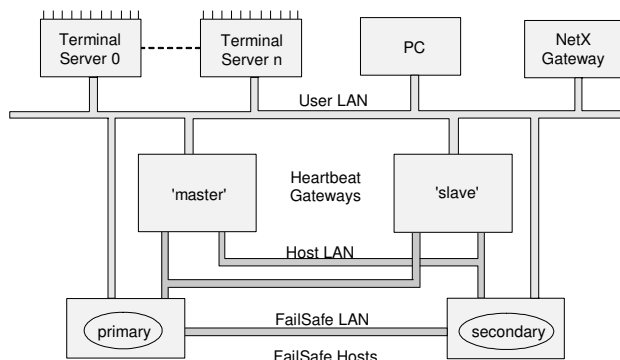
Reality Heartbeat is an optional software feature for selected FailSafe platforms providing enhanced resilience and usability.

Heartbeat enhances resilience by automatically detecting primary system failure. In this event the secondary system is re-configured as the primary and users are switched to it. Users are then automatically logged on to the top level of their current application so that they can continue working with minimal disruption to service.

Heartbeat provides a transparent connection to the primary database. A user connects to the database without needing to know the host machine names or the current primary.

System Architecture

A Heartbeat system consists of two UNIX host computers operating in FailSafe mode and connected to the user network via Heartbeat switching software located in two or more Heartbeat gateway computers.



FailSafe Configuration

The FailSafe configuration comprises two systems operating in 'hot standby' configuration. One machine, the 'primary', is nominated to carry out the processing and printing for users. The other machine, the 'secondary', acts as the standby in case the primary fails.

Heartbeat Gateways

Transaction Logging replicates database updates from the primary onto the secondary via a dedicated FailSafe LAN. For additional security, for audit trails and to provide for database recovery, updates are also recorded in disk resident 'clean log' files. For more details of Transaction Logging and FailSafe, refer to Product Datasheet *Reality Resilience*.

Heartbeat gateways are computers that run switching software to provide an intelligent interface between users and the FailSafe system. Each gateway automatically connects a user logging on via the gateway to the primary database.

The gateway currently designated as the master also monitors the FailSafe hosts for a failure and initiates appropriate action to maintain user connections.

Multiple gateways need to be configured depending on the required user loading. At least two are necessary to ensure resilience if a gateway fails. Additional gateways provide increased throughput and higher levels of redundant resilience.

LAN Configuration

Four local area networks (LANs) interconnect a Heartbeat system:

- A User LAN, which connects users to the Heartbeat gateways and FailSafe hosts.
- Two Host LANs, which connect the Heartbeat gateways to the two FailSafe hosts.
- FailSafe LAN, which provides dedicated logging to the secondary.

The gateway LAN configuration ensures resilience in the event of various forms of LAN failure.

Heartbeat FailSafe Feature

System Operation

Users requiring automatic switchover access the primary database via a Heartbeat gateway. Other users can log on directly to a host system via the user LAN.

Users logging on to a Heartbeat gateway are automatically connected to the primary database. A standby connection to the secondary is built for use in the event of a switchover. The standby connection is 'stalled' before the application is executed. The stall is released in the event of a switchover.

System Health Monitoring

Heartbeat monitors the FailSafe system using a Monitor process on each host and an Arbitrator process on the master gateway.

The Monitor comprises UNIX daemons that continually check the system hardware, UNIX operating system and Reality environments and send regular Heartbeat pulses to the Arbitrator to indicate correct functioning.

If the Monitor detects a fatal error, it reports the failure on the system console and sends a message to the Arbitrator with the reason for the failure.

The Arbitrator monitors Heartbeat pulses from the hosts. If it fails to receive a number of consecutive pulses from a Monitor (or it receives a 'fail' message from a Monitor) it deems the system to have failed. It then initiates appropriate recovery action.

Periodically, each Monitor sends 'request status' messages to the Arbitrator and waits for a reply to check that it is running. If, after a number of requests, no reply is received, the Arbitrator is deemed to have failed and the failure is reported to the console.

Failure and Recovery

Host System Failure:

If the Arbitrator detects a fatal error on one of the hosts, it reports the failure to the Heartbeat master gateway console and, if possible, to the console connected to the failed machine. Heartbeat then initiates appropriate recovery actions.

If the primary fails, Heartbeat users are automatically switched to the secondary on the standby machine where they are returned to the logon Proc that can return them to the main menu of their current application with minimal disruption.

If the secondary fails, users continue to work without disruption while the secondary system is repaired.

When a failed system has been repaired, its database(s) must be recovered from the live system using standard Transaction Logging/FailSafe procedures. Recovery can take place at any convenient time and does not disrupt service to users connected to the live system. When the databases on the two systems are resynchronized, FailSafe operation is re-established automatically with the repaired system assuming the role of secondary. The Heartbeat gateways automatically build standby connections for all users logged on. Thus, if a primary failure occurs at a later date, a standard switchover can take place.

Heartbeat Gateway Failure:

Failure of a Heartbeat Gateway causes all network connections via that gateway to time-out. Affected users must logon again via an alternative route that bypasses the failed component – alternative routing is usually configured so that this will be transparent to those users.

If the master gateway fails, another gateway must be re-configured as the master and the failed gateway repaired and reconfigured as a slave. If a slave gateway fails the Heartbeat system is unaffected apart from users routed through that gateway. The failed gateway is repaired and configured again as a slave.

Gateway Load Sharing

Load sharing software manages the assignment of new user connections so as to distribute the user load evenly across multiple gateways and prevent overload.

If the total user load on a Heartbeat system becomes abnormally high this software directs new logons away from an overloaded gateway towards more lightly loaded ones.

If a gateway fails, it is marked as unavailable for new connections and the load sharing software re-routes further logons via an alternative gateway.

Load sharing is supported for both OSI and TCP/IP protocols, and user-specific features of the external networking equipment.

Administration

A UNIX-based utility hbmenu displays menus and executes prompt-driven procedures to provide a user-friendly interface for administering a Heartbeat system.