



Low Latency Real-Time Computing on Multiprocessor Systems Running Standard Linux

Dimitri Sivanich, SGI
sivanich@sgi.com
1140 E. Arques Ave.
Sunnyvale, CA 94085

Low Latency Real-Time Computing on Multiprocessor Systems Running Standard Linux

- Through proper configuration and application of real-time strategies, a standard Linux kernel can run many applications requiring low latency and fast interrupt response without the need for an underlying real-time kernel.
- Latency Effects
 - Interrupt Response - time to process an interrupt and wake a sleeping thread
 - Jitter - periods of cpu unavailability during thread execution



Steps to Achieving a Low Latency Linux OS Installation

- Install a distribution that is suitable to real-time work in an isolated processor configuration.
 - Modify the kernel with fixes to unexpected latencies
or
 - Install a distribution already suited to real-time work (such as an SGI React supported distribution)
- Isolate a set of processors for real-time use:
 - Isolate from non-application processes and threads, load balancing effects, etc...
 - SGI React provides a configuration tool to easily accomplish this
- Design and run your application using recommended strategies
 - Pin to isolated cpu, lock memory, use shared memory for IPC, etc.

Highest Interrupt Response Times

Highest interrupt response times (in microseconds)		
4 CPU system	without REACT	with REACT
no-load	349.30	11.95
load	556.95	11.15
64 CPU system	without REACT	with REACT
no-load	484.95	10.50
load	545.35	18.75