

A dark, semi-transparent world map is centered in the background of the slide. The map shows the continents in a slightly lighter shade than the surrounding dark background.

# libtcr

Making the most of your cores  
Philipp Reisner  
Linuxkongress 2010, Nuremberg

We got cores  
We will get more!

Non-uniform memory access



Cache coherency is expensive,  
... will get more expensive!

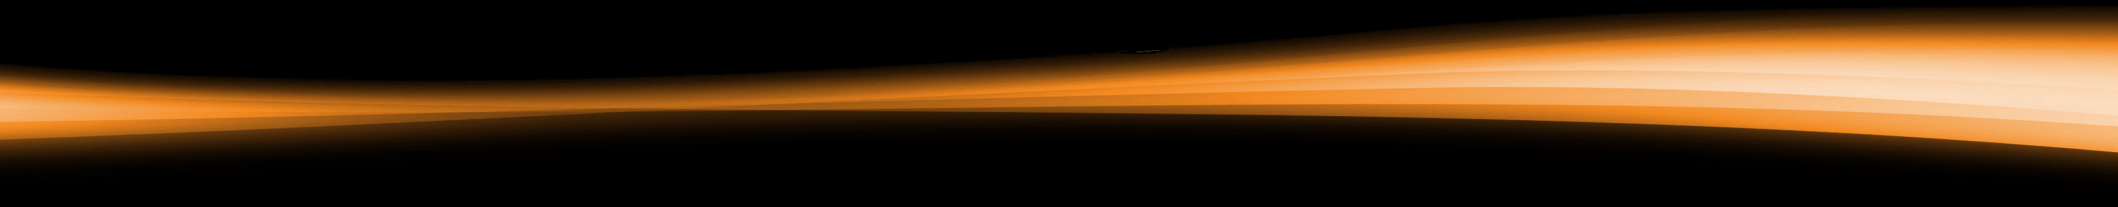


Server programs,

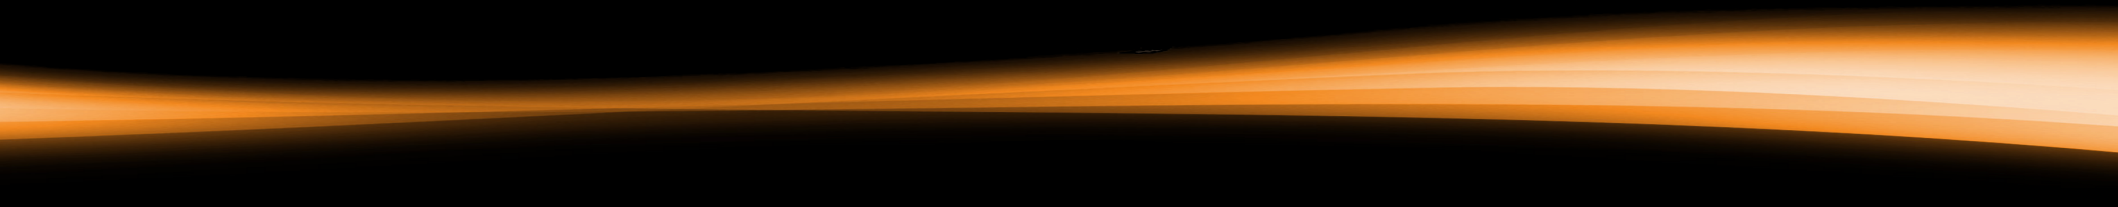
Server programs,  
Multi process - ... threaded - state machines



Server programs,  
Multi process - ... threaded - state machines  
complexity,  
low - low - hard

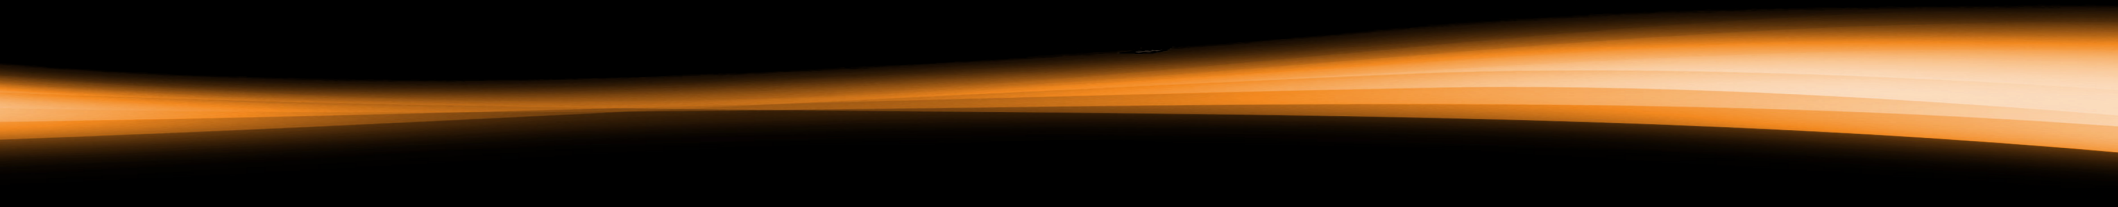


Server programs,  
Multi process - ... threaded - state machines  
performance on UP,  
low - better - best





Server programs,  
Multi process - ... threaded - state machines  
throughput scales with cores,  
yes - yes - no



The winner is...

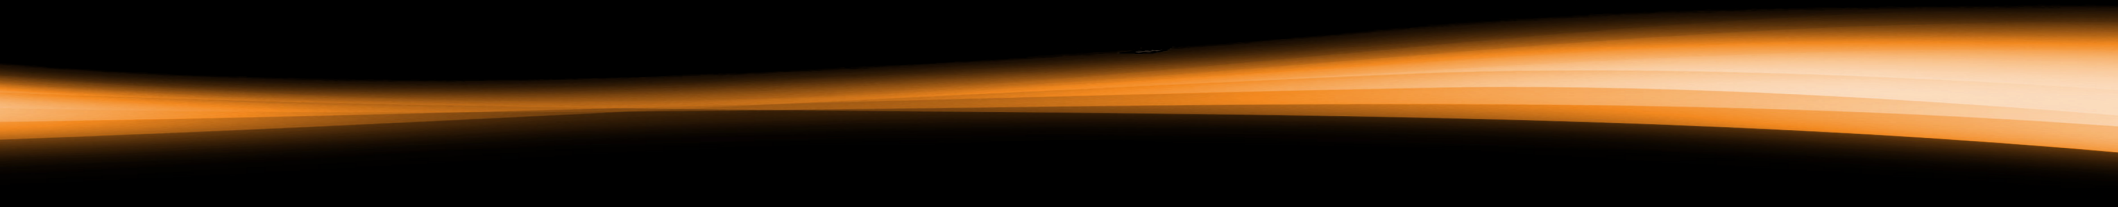


<b>complexity</b>	<b>+</b>	<b>+</b>	<b>-</b>
<b>uni processor</b>	<b>-</b>	<b>-</b>	<b>+</b>
<b>scales cores</b>	<b>+</b>	<b>+</b>	<b>-</b>
<b>cache locality</b>	<b>-</b>	<b>-</b>	<b>+</b>
	<b>multi proces</b>	<b>multi threaded</b>	<b>state machine</b>

A combination



program's state = what to do next



state machine = store state explicitly

threading = state implicitly on the stack

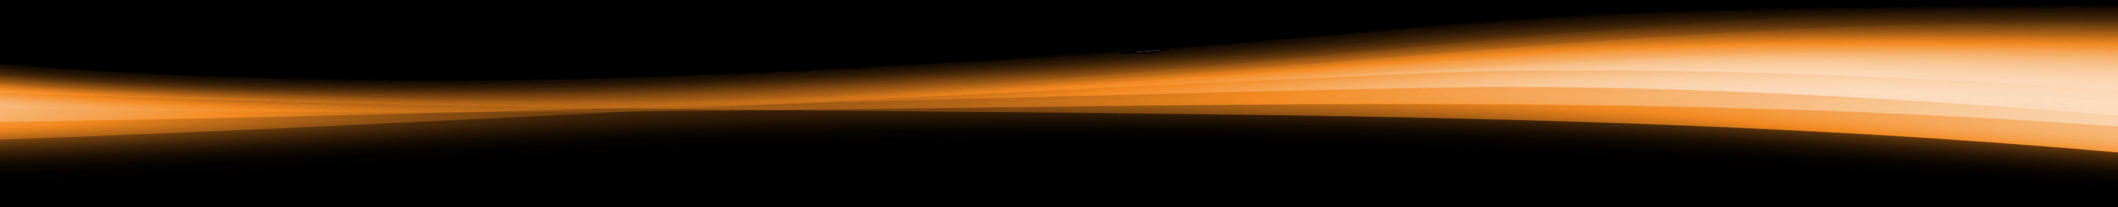


Switching stack = one assignment



libtrc is ...

...scheduling (stack switching)  
in user space  
on multiple kernel threads  
in parallel



N:M threading



Pipelining...

... with a thread per stage,  
is worst case for cache coherency

Process data on one CPU,  
read, process and write



Features,



Features,  
threads, mutexes, wait queues...





Features,  
wait\_fd, rearm\_fd - 2 priorities

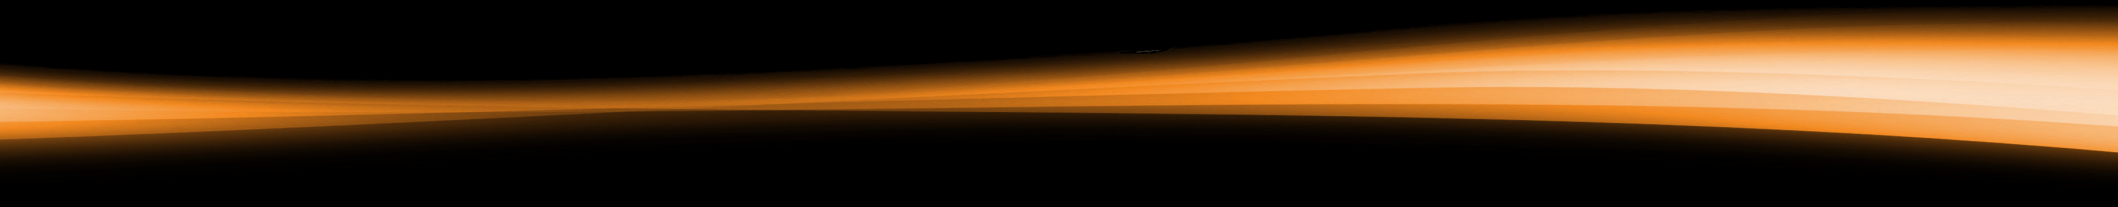


Features,  
tc\_signal

Features,  
parallel for, parallel statement macro



Event sources,  
FDs = FDs & timers, Unix signals, AIO



Internally,  
epoll

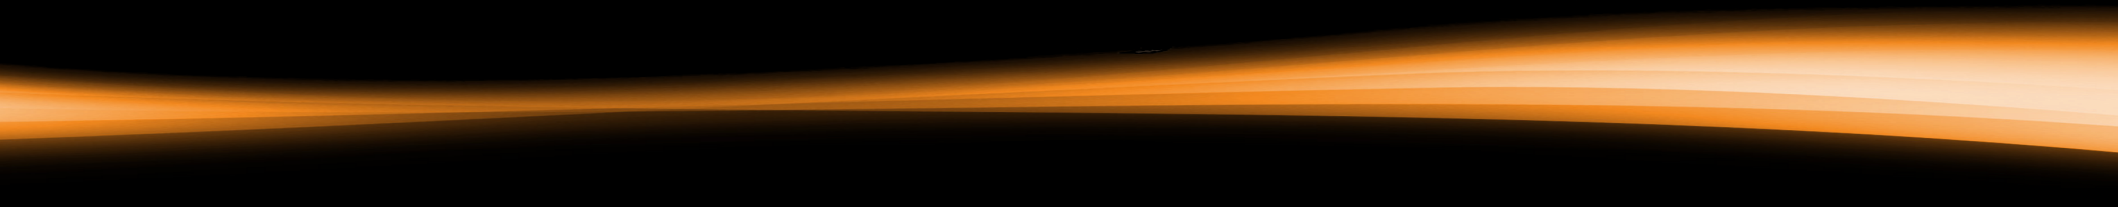
Internally,  
epoll and a pthread per core



Requires,  
Linux 2.6.25




Area to improve,  
combination of  
cpu intense and  
IO multiplexing workloads





# Drawbacks, Debugging aids



Need more information?

<http://oss.linbit.com/libtcr>

<http://git.drbd.org/libtcr.git>

The LINBIT logo is located in the bottom right corner of the slide. It features the word "LINBIT" in a bold, white, sans-serif font. The letter "I" is stylized with a vertical line through it. The logo is set against a dark background with a glowing orange and yellow gradient at the bottom of the slide.