Paul E. McKenney, IBM Distinguished Engineer, Linux Technology Center
Member, IBM Academy of Technology
Linux Plumbers Conference, New Orleans, LA, USA September 18, 2013

# But What About Updates?

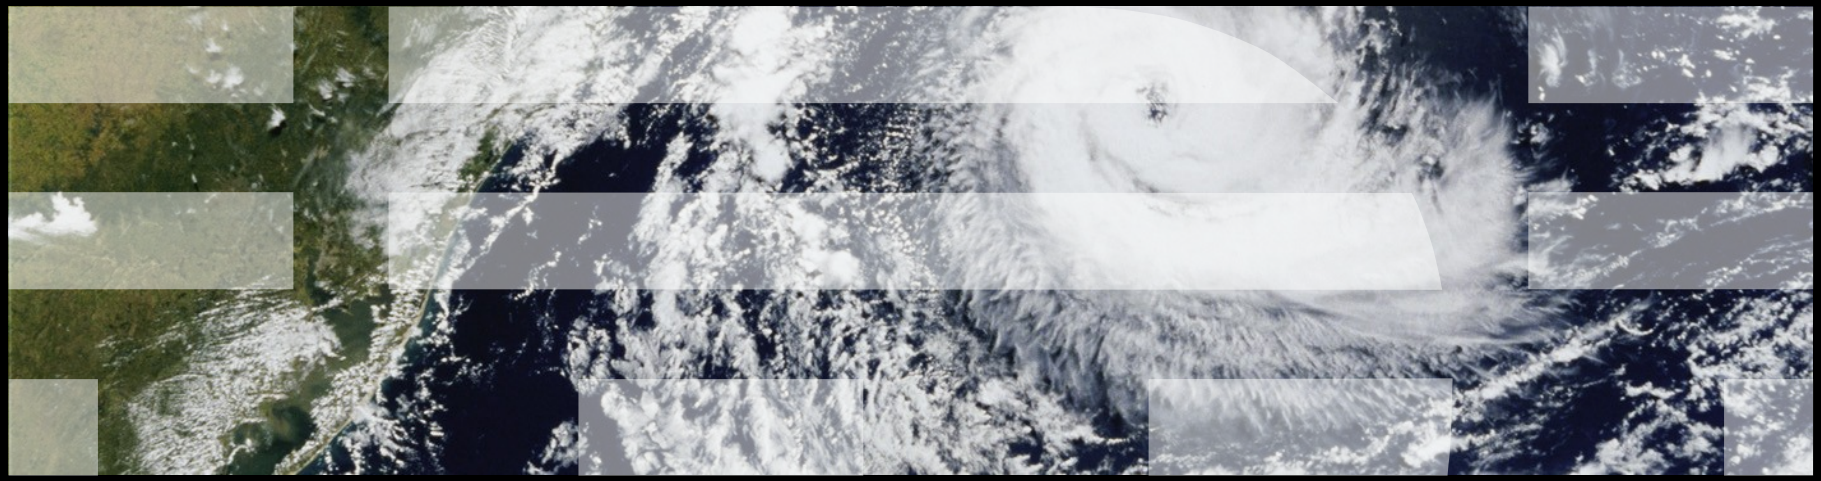

# Read-Mostly Workloads Scale Well

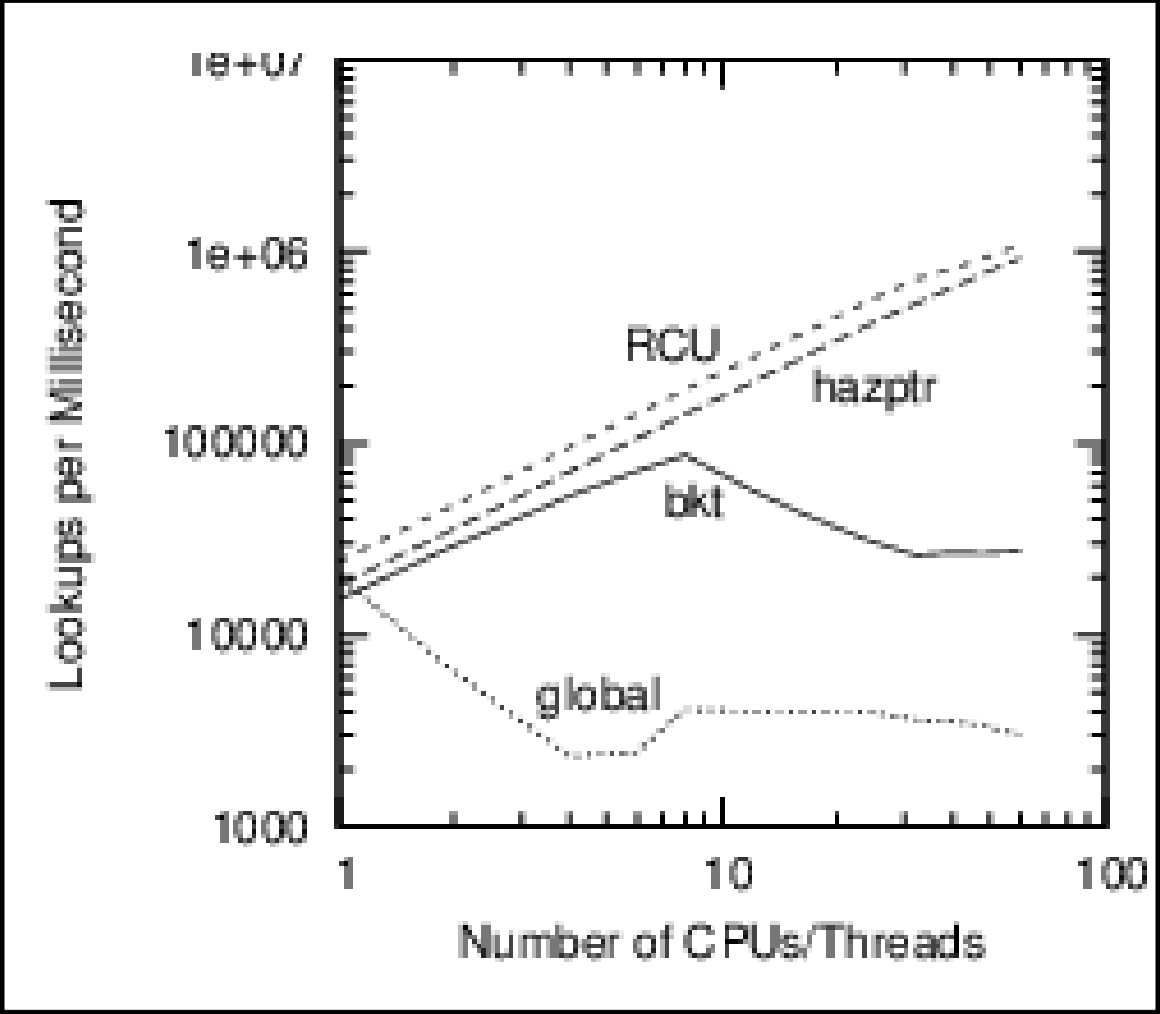

# Update-Heavy Workloads, Not So Much...

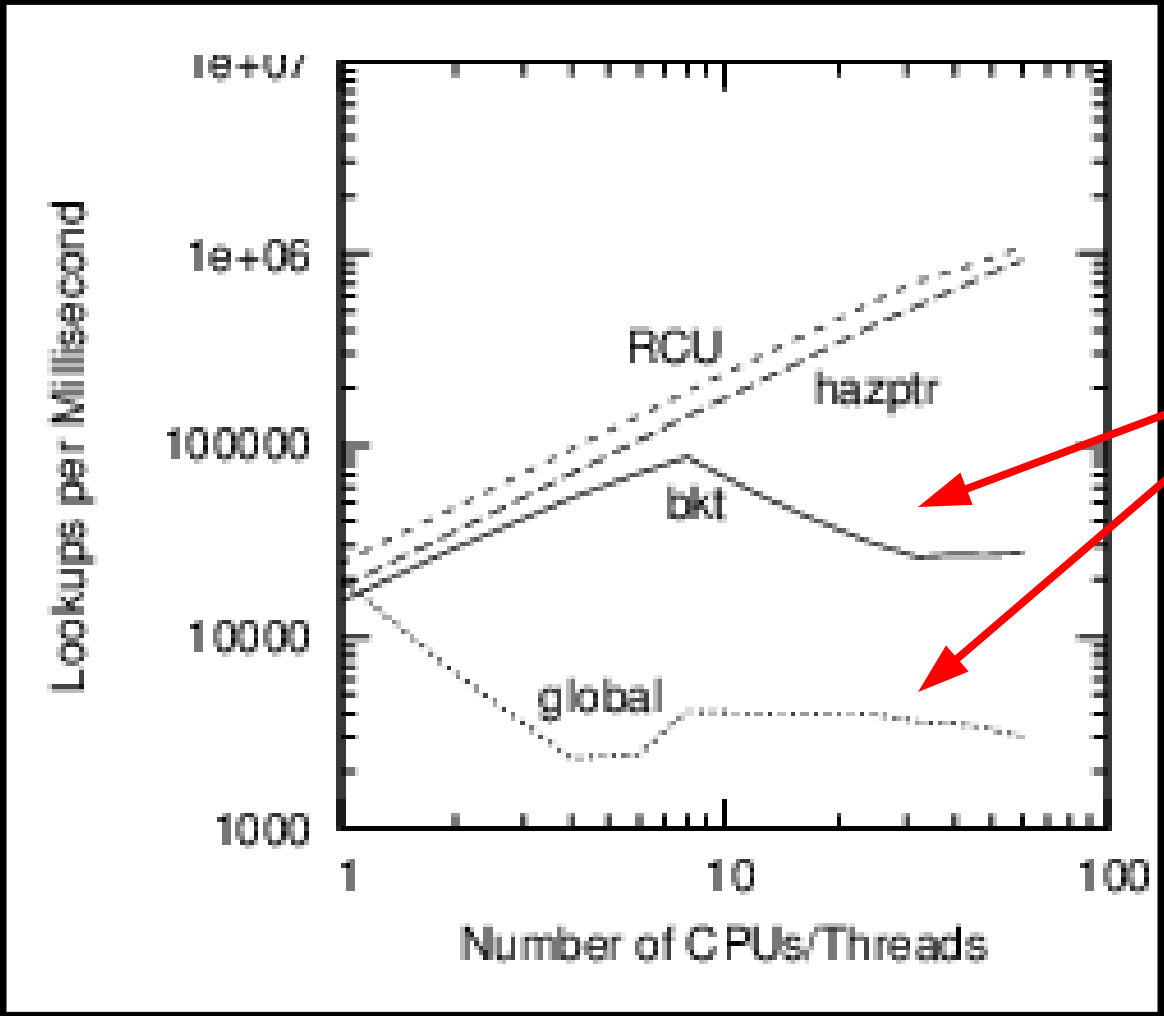

# But There Are Some Special Cases

# But There Are Some Special Cases

- Split counters (used for decades)
  - Have a per-CPU/thread counter
  - For updates, each CPU/thread non-atomically updates its own counter
  - For reads, sum all the counter
  - Rely on commutative and associative laws of addition
  - Plus rely on short-term inaccuracy permitted for statistical counters
  - Constant work done for updates, linear scaling, great performance
- Per-CPU/thread processing (perfect partitioning)
  - Huge number of examples, including the per-thread/CPU stack
  - But not everything can be perfectly partitioned

# Special Case: Stream-Based Applications

- Adrian Sutton of LMAX presented this at linux.conf.au 2013:
  - http://www.youtube.com/watch?v=UvE389P6Er4
  - http://lca2013.linux.org.au/schedule/30168/view_talk
- Only two threads permitted to access a given location
- Use fixed-array circular FIFOs to pipe data between data-processing stages (represented by individual threads/CPUs)
- Get nearly uniprocessor performance, especially for heavy-weight processing

# Example Stream-Based Application

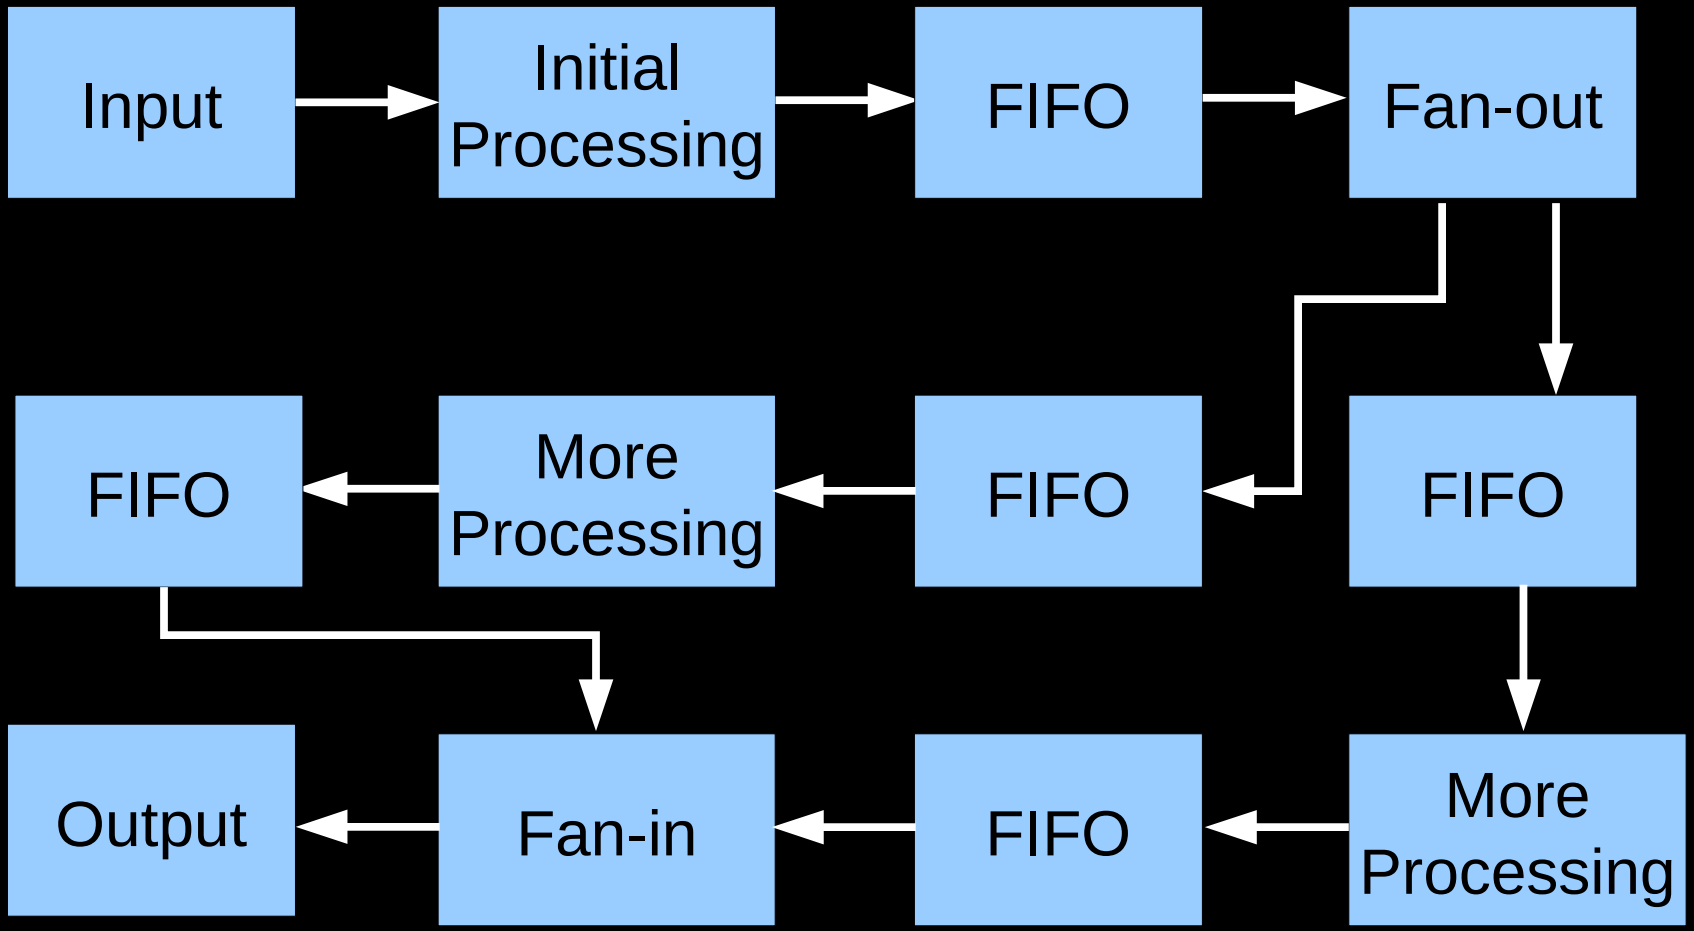

# Other Approaches

- Hardware transactional memory
  - You saw Andi Kleen's talk
- More sophisticated uses of associativity and commutattivity
  - Research topic, some progress being made
  - And they are using the Linux kernel as a test case!
- Your ideas here!!!

# Summary

- We are farther along with read-mostly methods than with update-heavy methods
- But there are some good approaches for update-heavy workloads for some special cases
  - Split counters
  - Stream-based applications
  - Hardware transactional memory
  - Maybe some more

# Legal Statement

- This work represents the view of the author and does not necessarily represent the view of IBM.
- IBM and IBM (logo) are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries.
- Linux is a registered trademark of Linus Torvalds.
- Other company, product, and service names may be trademarks or service marks of others.

# Questions?