

Acunu Storage Platform

Richard Low
rlow@acunu.com

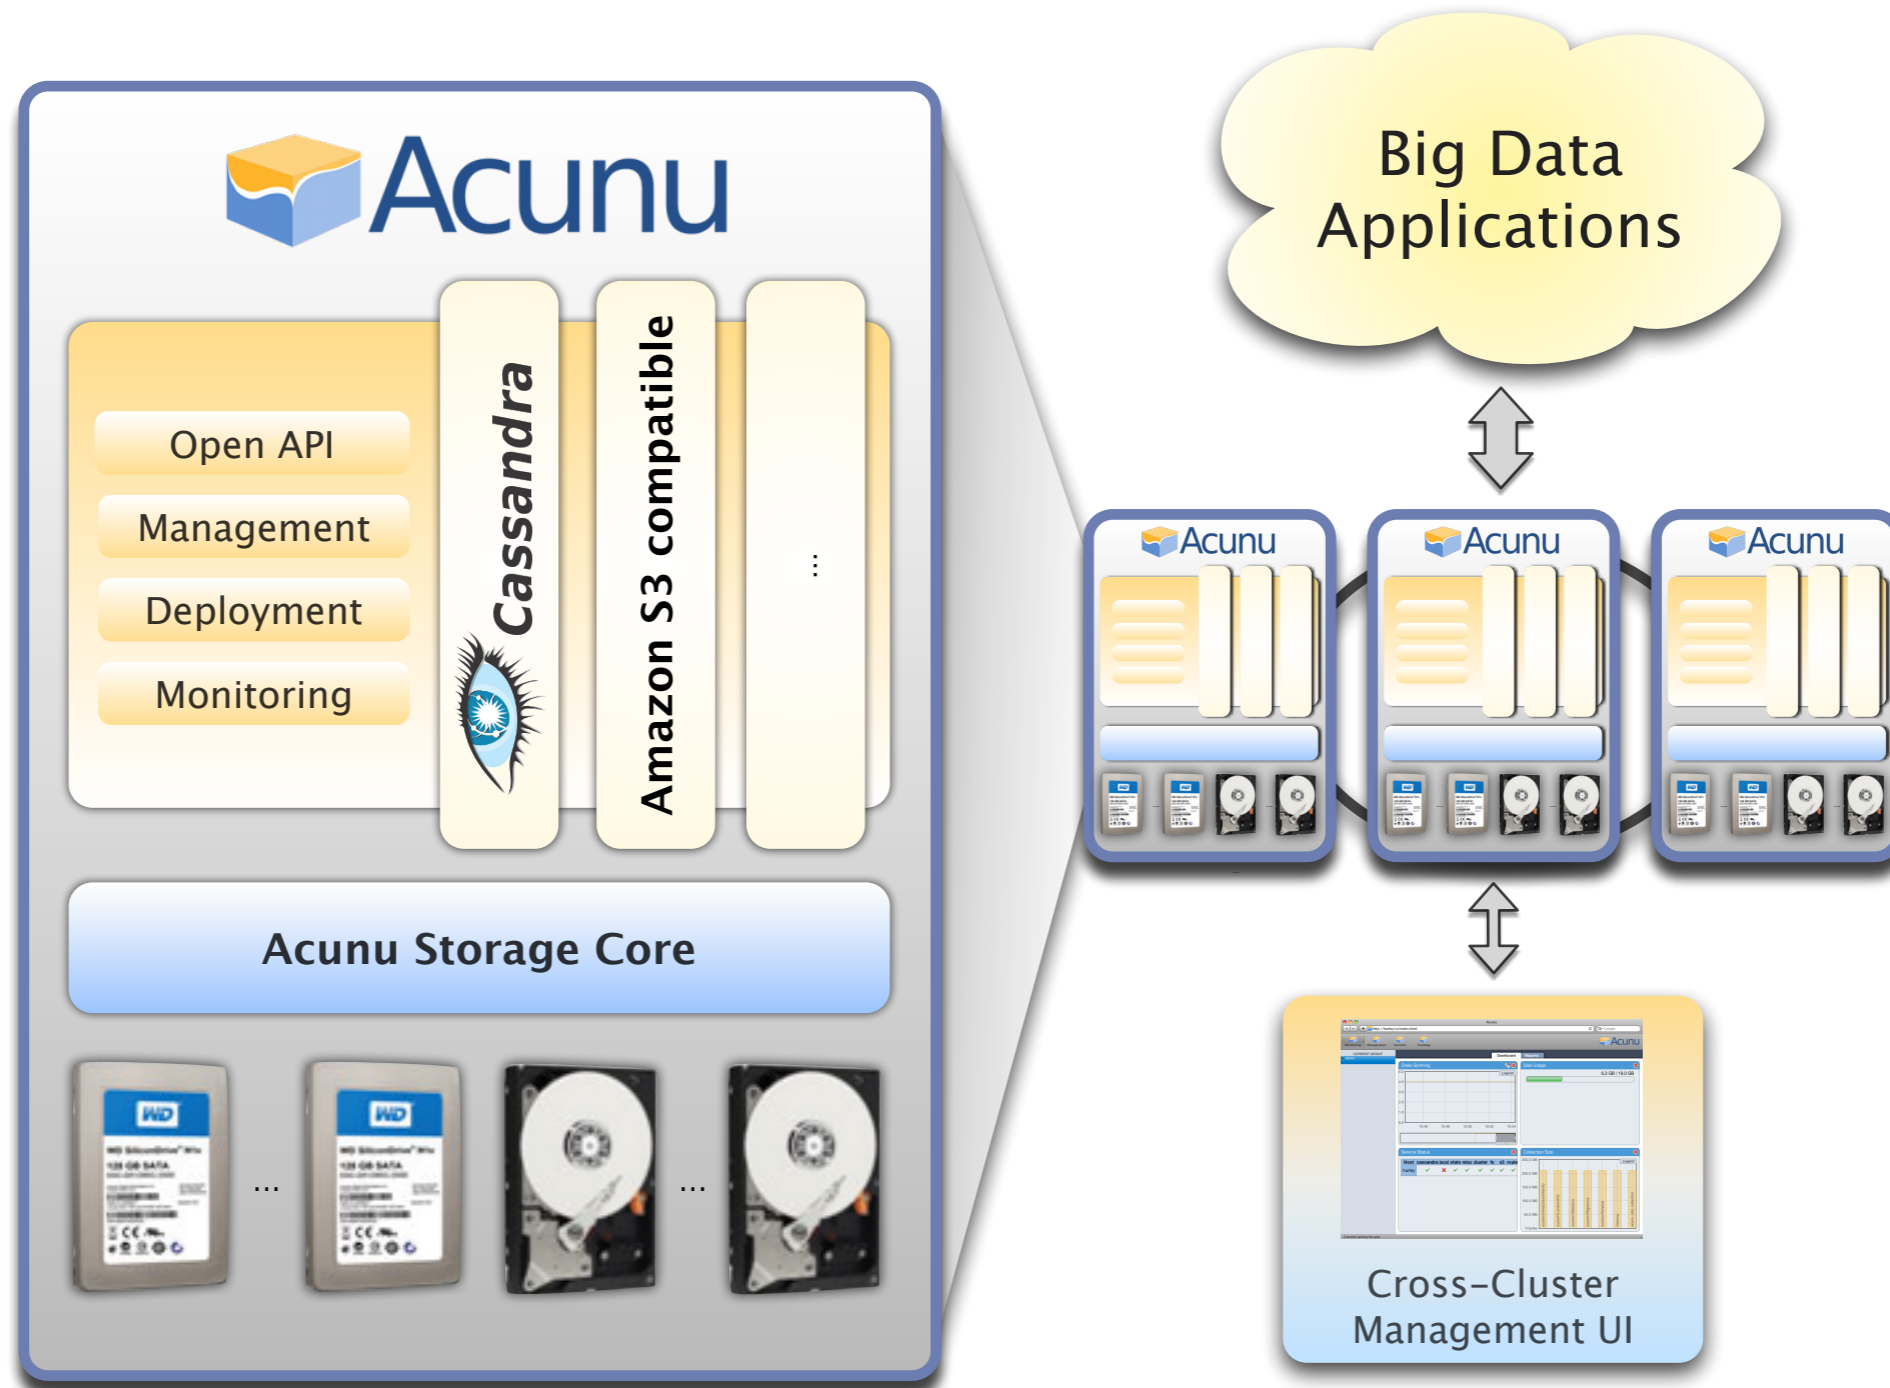


Outline

- What is the Acunu Storage Platform?
 - Overview
 - Features
 - API
- What does it do for Cassandra?
 - Predictable performance
 - Fast range queries

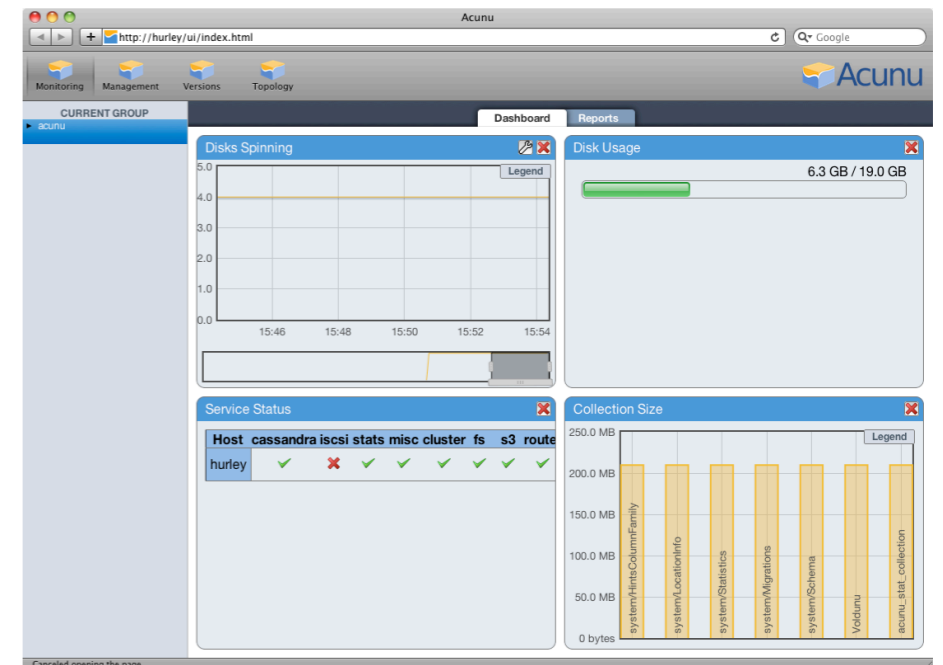


Overview



Acunu Management Framework

- Web UI
- Cluster management and monitoring
- Open APIs



Acunu Storage Core

- Open source storage stack
- Linux kernel module
- key-value interface



Acunu Storage Core (2)

- Add high performance to any big data application
- Easily use full disk bandwidth of HDDs and SSDs
- Fast disk rebuild (no more RAID)
- Versioning

key-value interface

`get(key)`

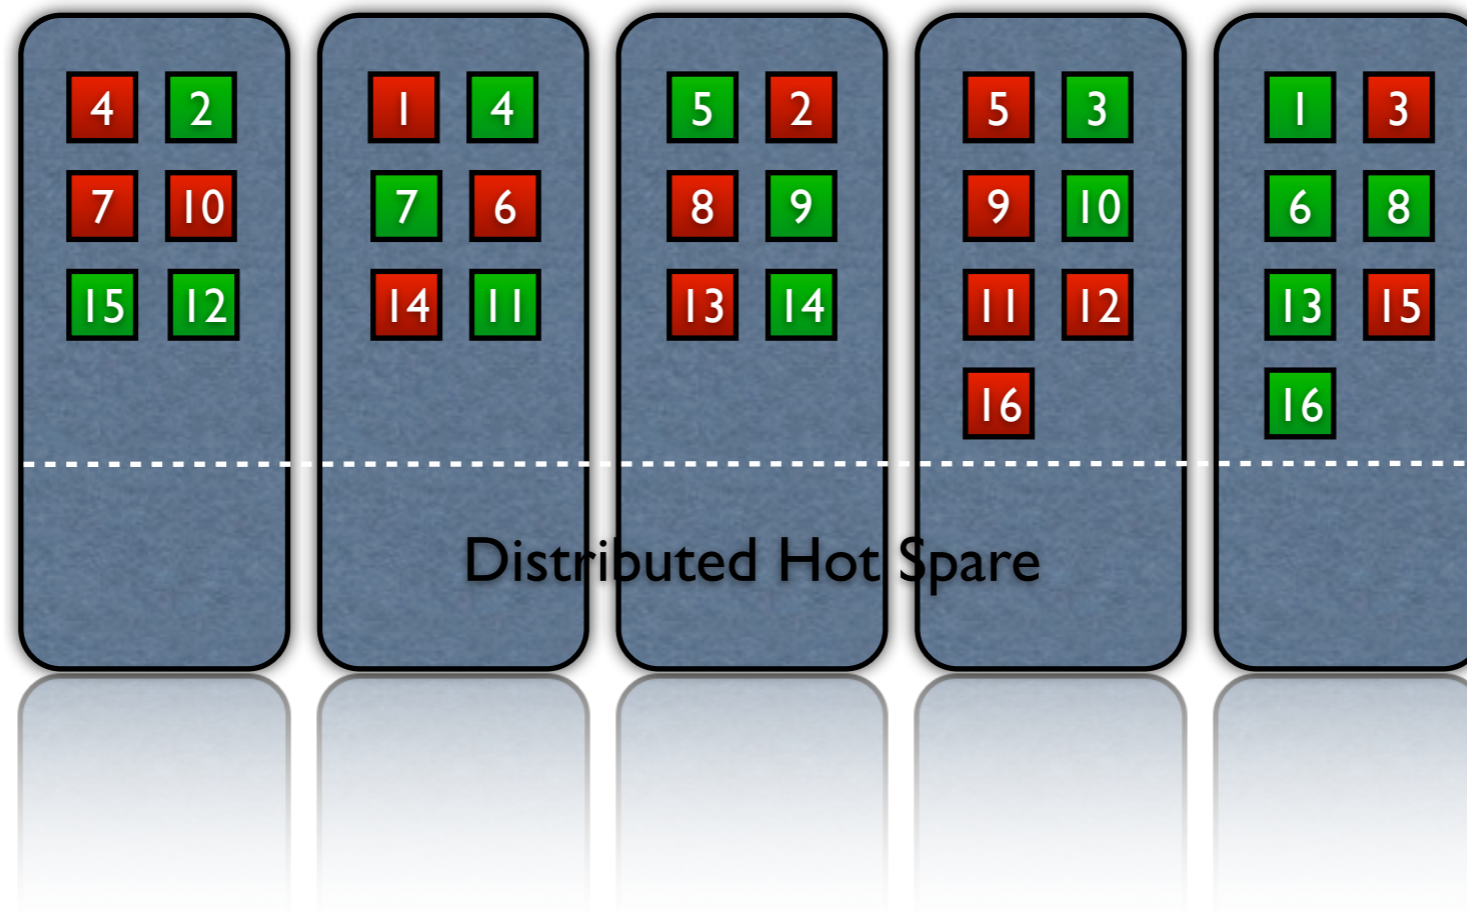
`put(key, value)`

`iterator(start, end)`

- Multi-dimensional keys (e.g. cassandra row, column)
- Very fast shared-memory interface

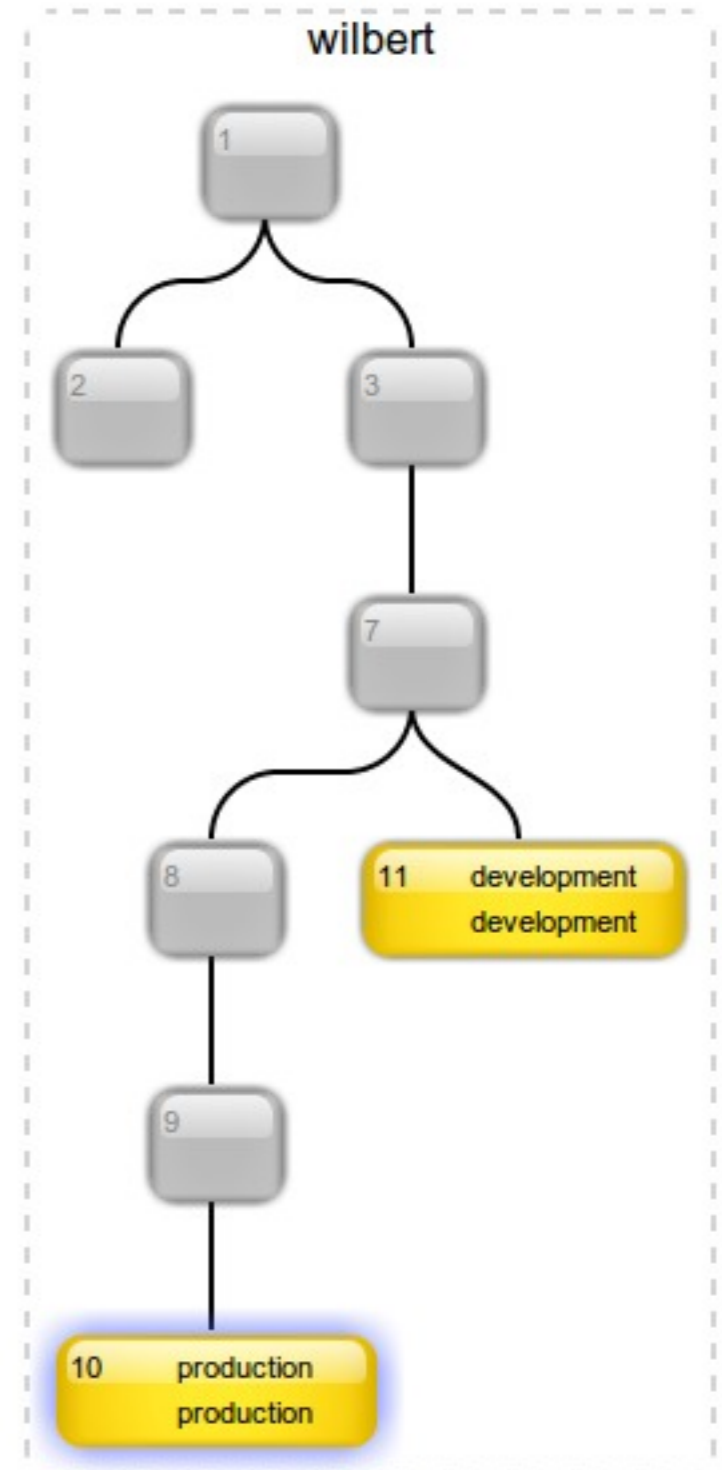
Fast disk rebuild

- Rebuild time depends on #disks, not size
- Bandwidth of hot spares not wasted

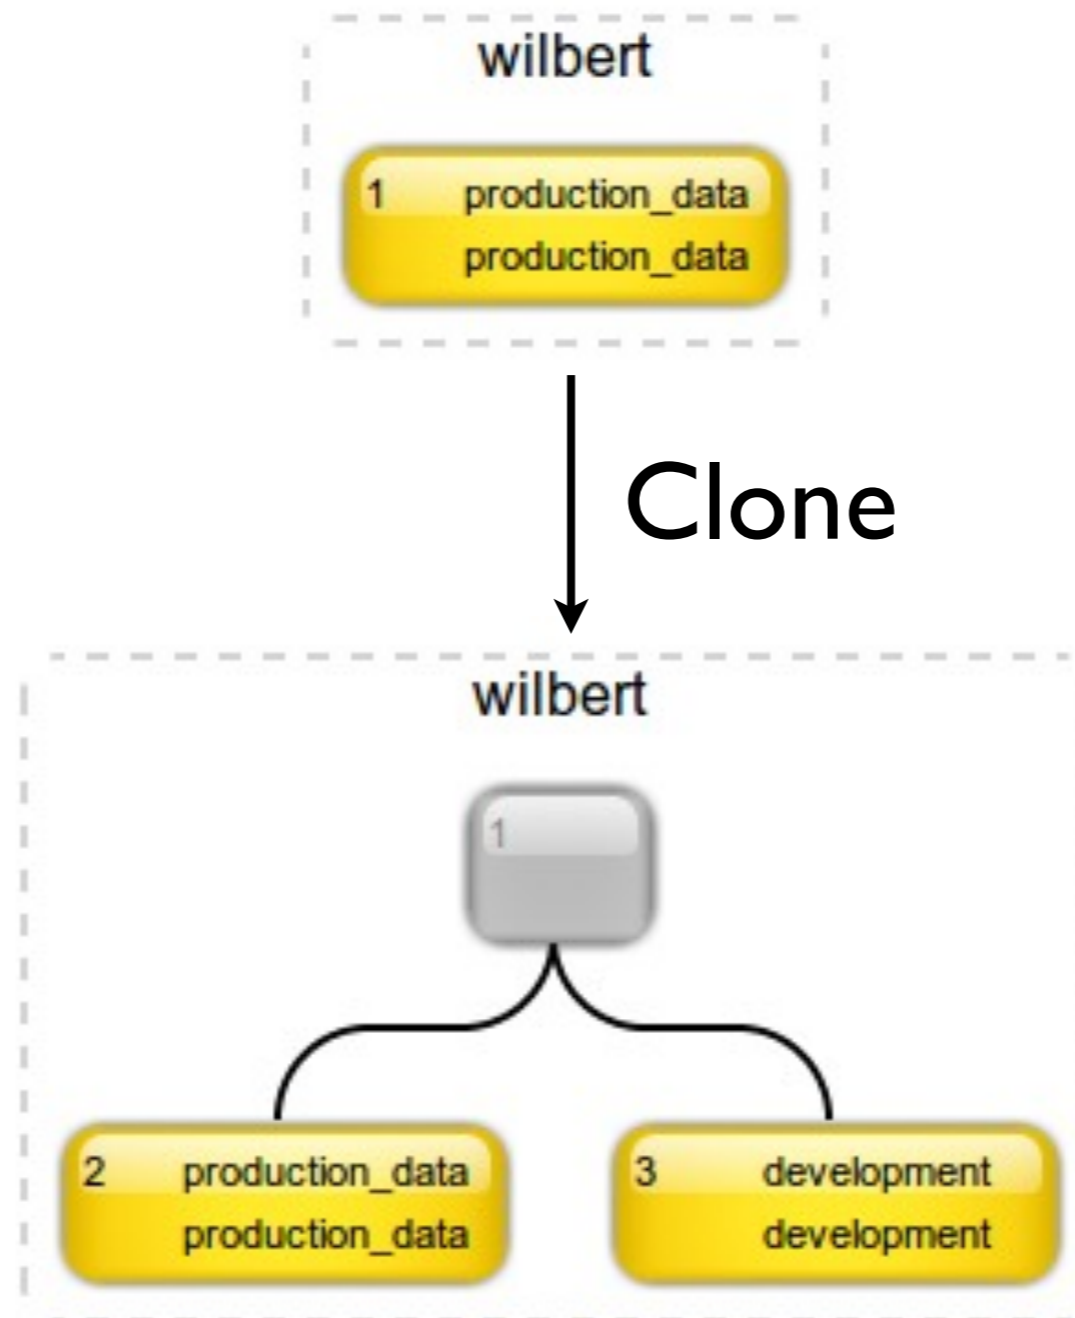


Versioning

- Instantaneous snapshots
- Avoid copying data
- Maintain efficient reads
- Obvious use: backup




Versioning: development clone



Versioning: historical analytics

- Keep snapshots at important times
- Run and rerun analytics at a later date



Big Data Applications Running on Acunu

- Cassandra



- S3



- Voldemort

- Memcache



- Hadoop




Cassandra adapter

- Replaces memtables, SSTables
- Replaces ColumnFamilyStore
- Same external interface

Insert rate

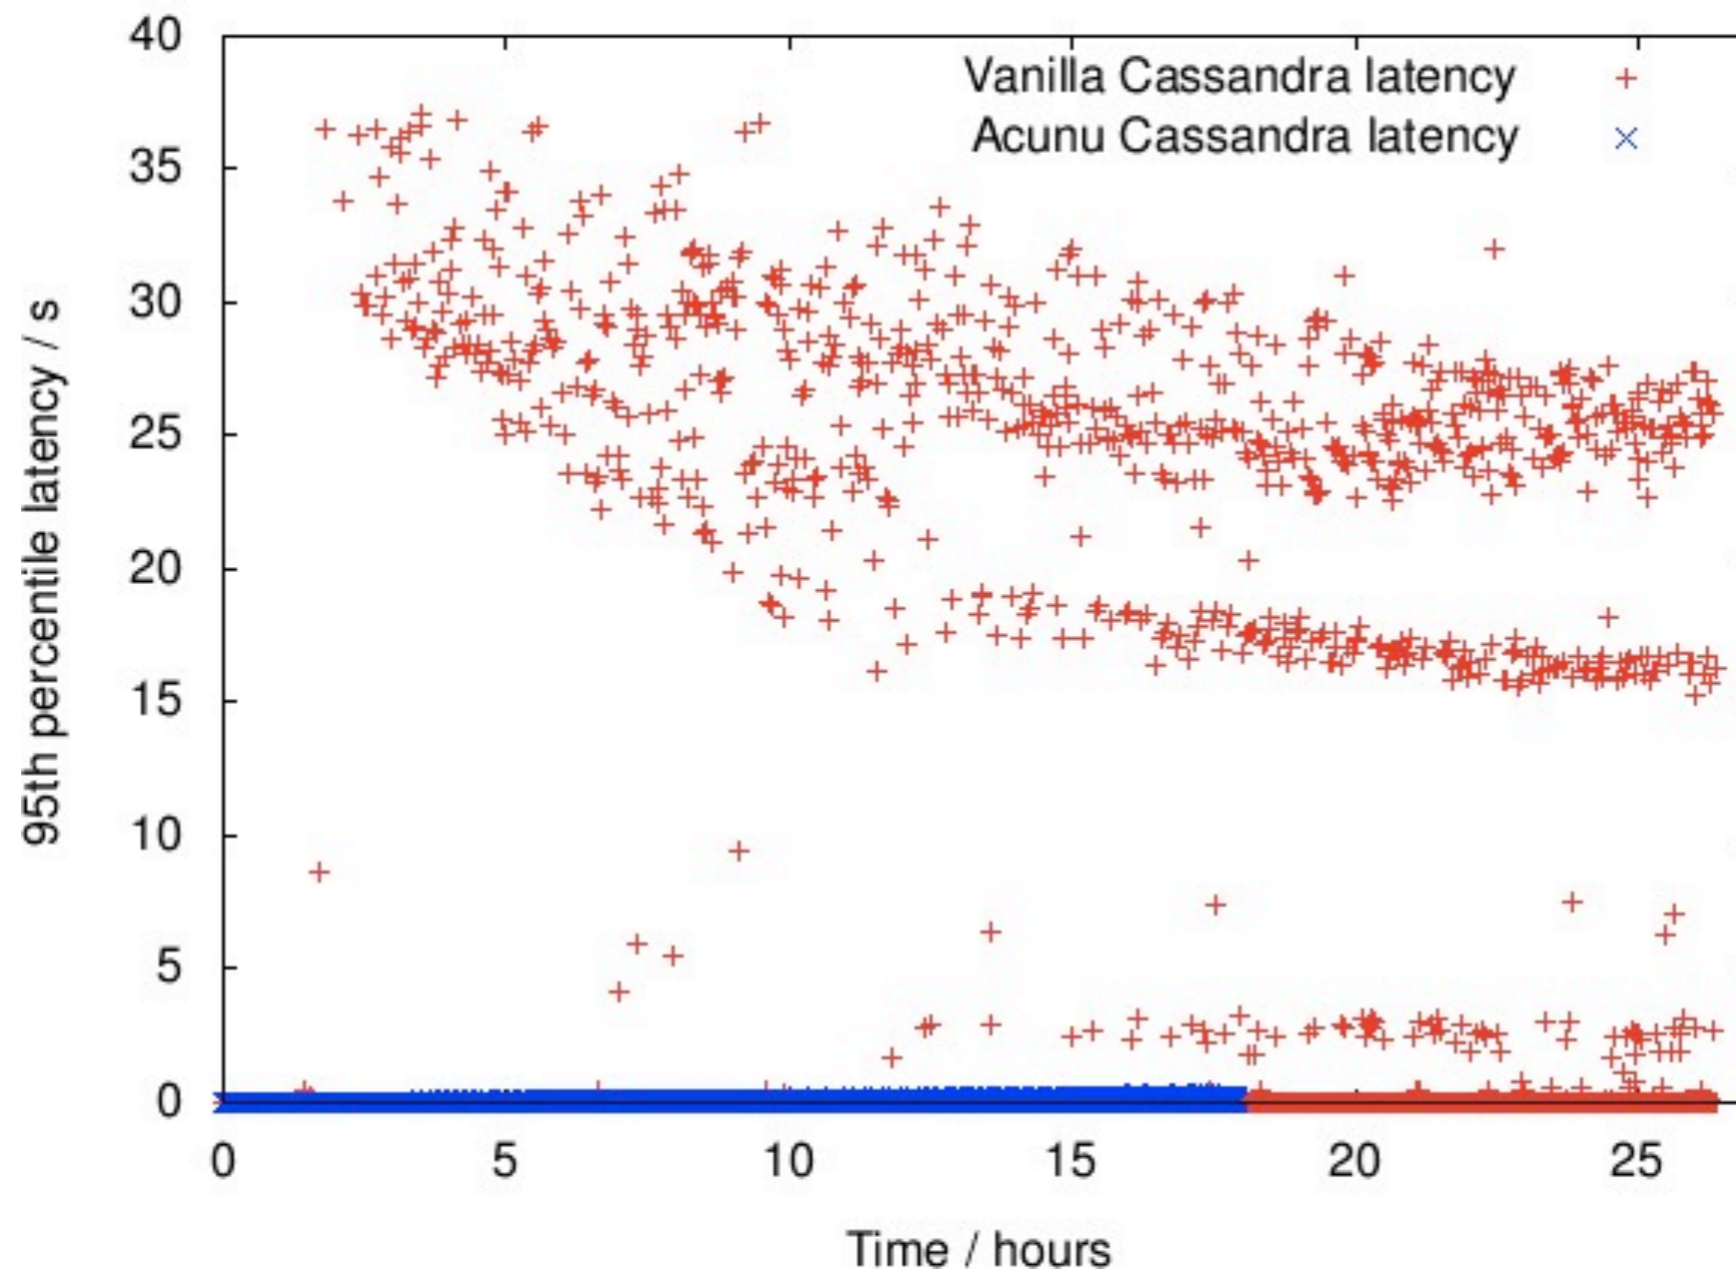


Insert latency

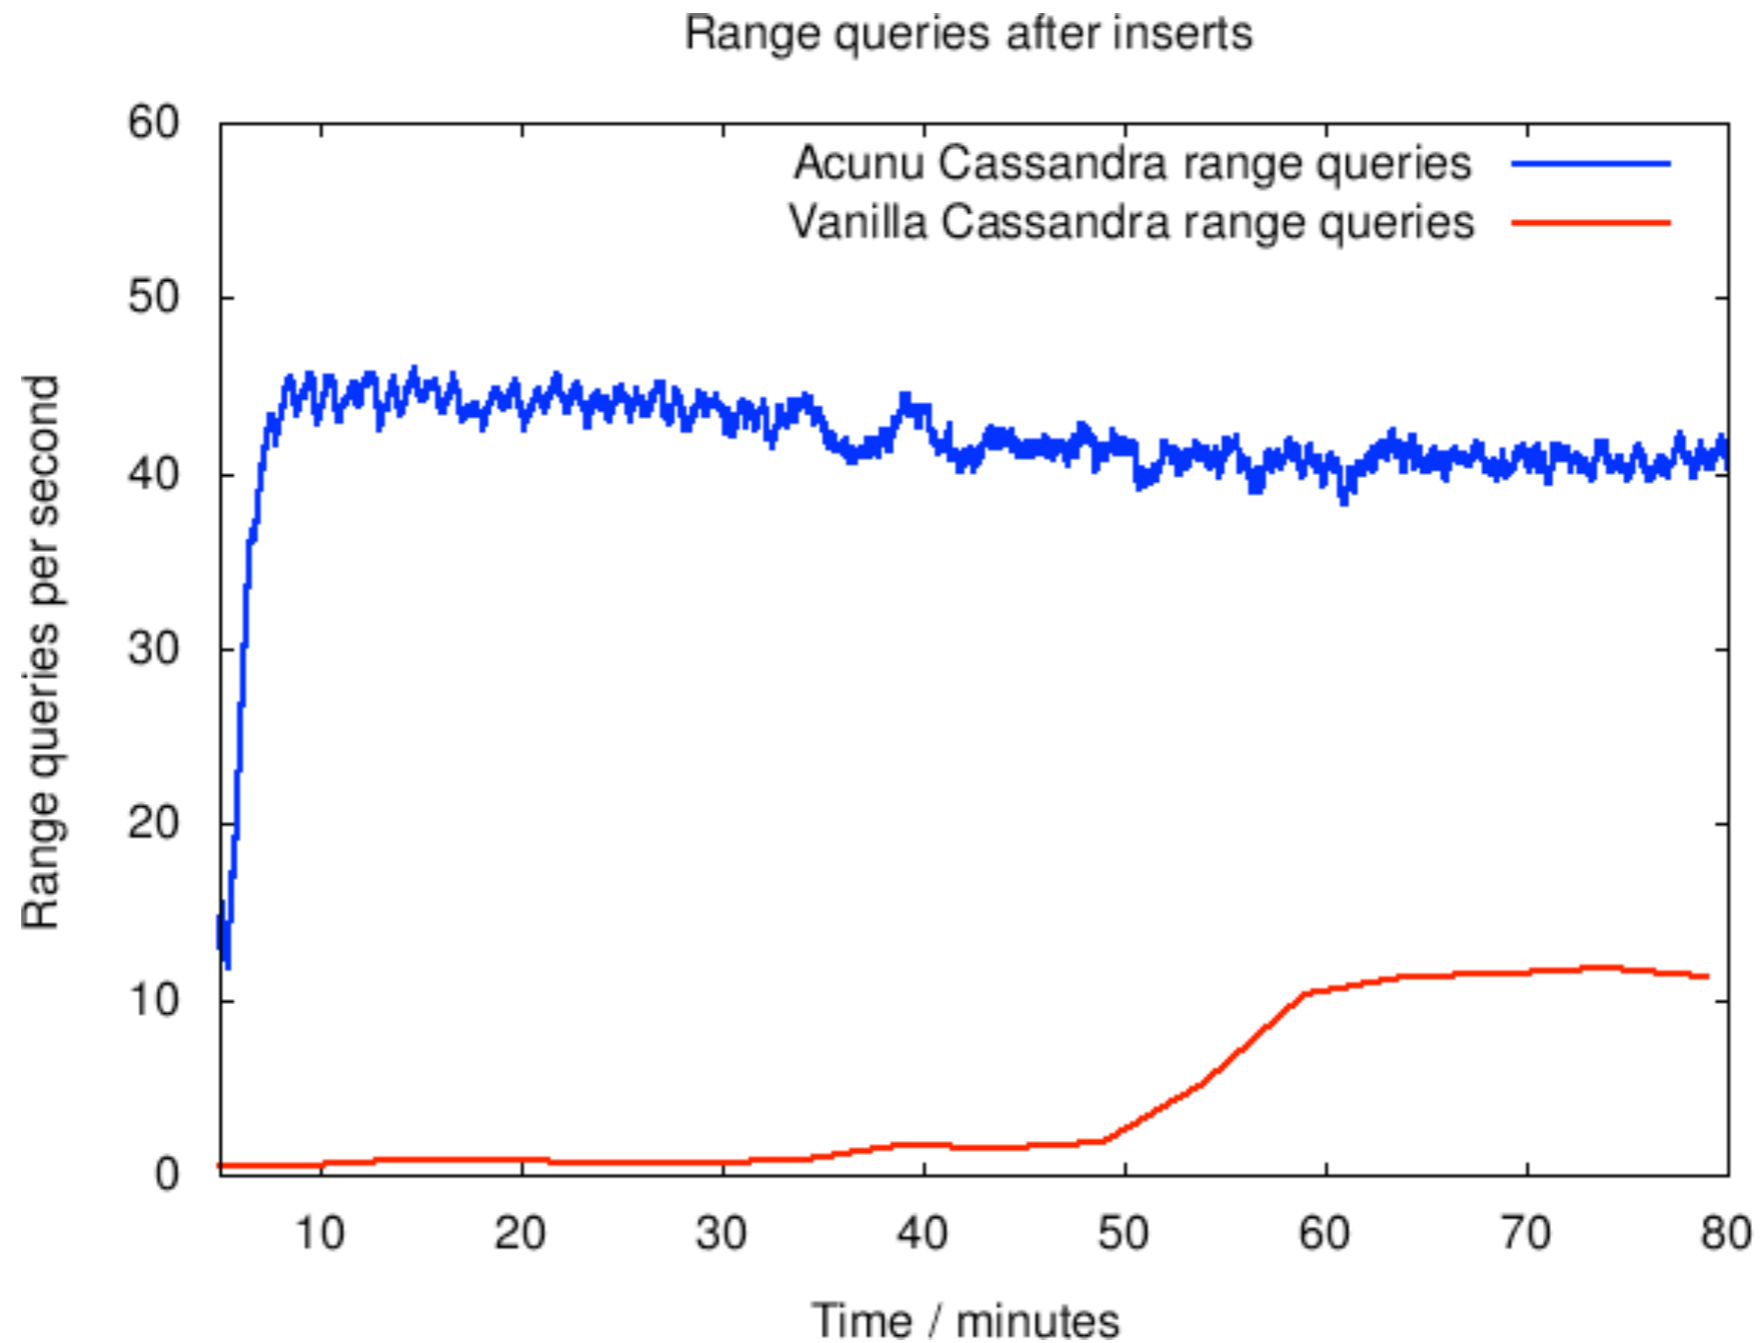


Latency

95th percentile insert latencies



Range queries



How?

- Carefully optimised C code
- No garbage collection
- Full merge deamortisation
- No major compactions
- Use full disk bandwidth
- No tuning required

How to use it

- Download open beta (www.acunu.com/beta)
- Get source code (bitbucket.org/acunu)
- Demo tomorrow



Download Beta 