

Crimson Project Update

Bill Scales, IBM




CPU Trend Data

- Clock speed plateaued
- More Cores



Concurrency is key

50 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

Goals:

- Lower CPU concurrency cost
- Utilize modern storage

How?

- “Classic” OSD replacement
- **Efficient** rewrite the I/O path



CPU Stats Comparison

	<u>Classic</u>	<u>Crimson</u>
Context switches / Sec	40,878	98
CPU migrations / Sec	6,524	None
Lock contention	455,000	400
Page faults / Sec	5	0.7
Cache Miss rate	12.86%	4.18%
CPUs utilized (out of 32)	24.2	31.78



Squid

- **Tech Preview** (Bluestore)
- RBD with Replication
- Automated Test Suite
- Snapshots
- Multicore Stability
- Recovery/Backfill
- Initial Scrub

Tentacle

- **SeaStore Tech Preview**
- Read Performance
- PG Splits
- Seamless deployment
- User Docs

Umbrella

- Included in default RPMs
- Cephadm / Seastore
- QoS / Mclock
- SeaStore laddr_t -> 128bits
- Erasure coding merged
- CephFS supported
- RGW supported

- PG Merging (pool shrink)
- Enable and verify fast EC
- Scrub scheduling
- SeaStore throughput improvements
 - Transaction conflict ratio
 - Onode/obc caching
 - Logical bucket cache

Join Crimson

- Backlog
 - <https://github.com/orgs/ceph/projects/8>
- Tracker
 - tracker.ceph.com/projects/crimson/issues
- Docs
 - docs.ceph.com/en/latest/dev/crimson
 - ceph.io/en/news/crimson
- Slack #crimson
- Weekly Meeting

