

Projects Since HPTS 2015

Charlie Johnson
(and interns and others)

HPTS '15



I'm going down!

**(How to recover from failure in
milliseconds)**

Charlie Johnson

Availability depends upon Time to Recover

- Mean Time Between Failures (MTBF) is important.
- Mean Time to Recover (MTR) is much more important.
- $\text{Availability} = A = \text{MTBF} / (\text{MTBF} + \text{MTR})$
- $\lim(\text{Availability}) \approx 1$
MTR $\rightarrow 0$
- As the Mean Time to Recover decreases, Availability approaches 100%
- \therefore Probability of Failure = $F = 1 - A \approx 0$
- \therefore Reliability = $1/F = 1/(1 - A) \approx \text{inf}$
- HPE Nonstop has implemented this and now does node takeovers in milliseconds: called “CPU Broadcast”: by hacking the NMI driver to send out a death multicast message. They now have the fastest takeover in the world.

TxHPC at NVMW 2017

PERSISTENT REGIONS THAT SURVIVE NVM MEDIA FAILURES

Onkar Patil, Mesut Kuscu, Tuan Tran, Charles Johnson, Joseph Tucek,
Harumi Kuno

Hewlett Packard Labs, Palo Alto, CA

NVMW 2017

What are the details?

Take Six: Software RAID 6 [10 shelves on separate Zbridges on one ZSwitch: 8 + 2 parity] at ~20% cost, using a "free band gap" (concurrency sized) and stripe iteration through the arrays, as if the data were moving up an escalator through the FAM (with much less overhead and work than Take Five.)

1. Stripes must be atomically written, but NOT update-in-place (the poison apple), so there needs to be a free band gap to accommodate concurrency. This implies that all array access needs to get on the stripe escalator at the appropriate point and "Mind The Gap".

FAM failures are recovered by reconstructing lost data or parity onto a spare shelf, or simply in DRAM for the demo

FAM Stripes

FAM Data Block

2. Using the page size of 4096 for the FAM Data Block, you get 512 64-bit floating or integer array elements, and then 4096 of said array elements per stripe for thread parallel data and parity computation in cores (one FPU and one Integer Unit per core)

FAM Parity Block

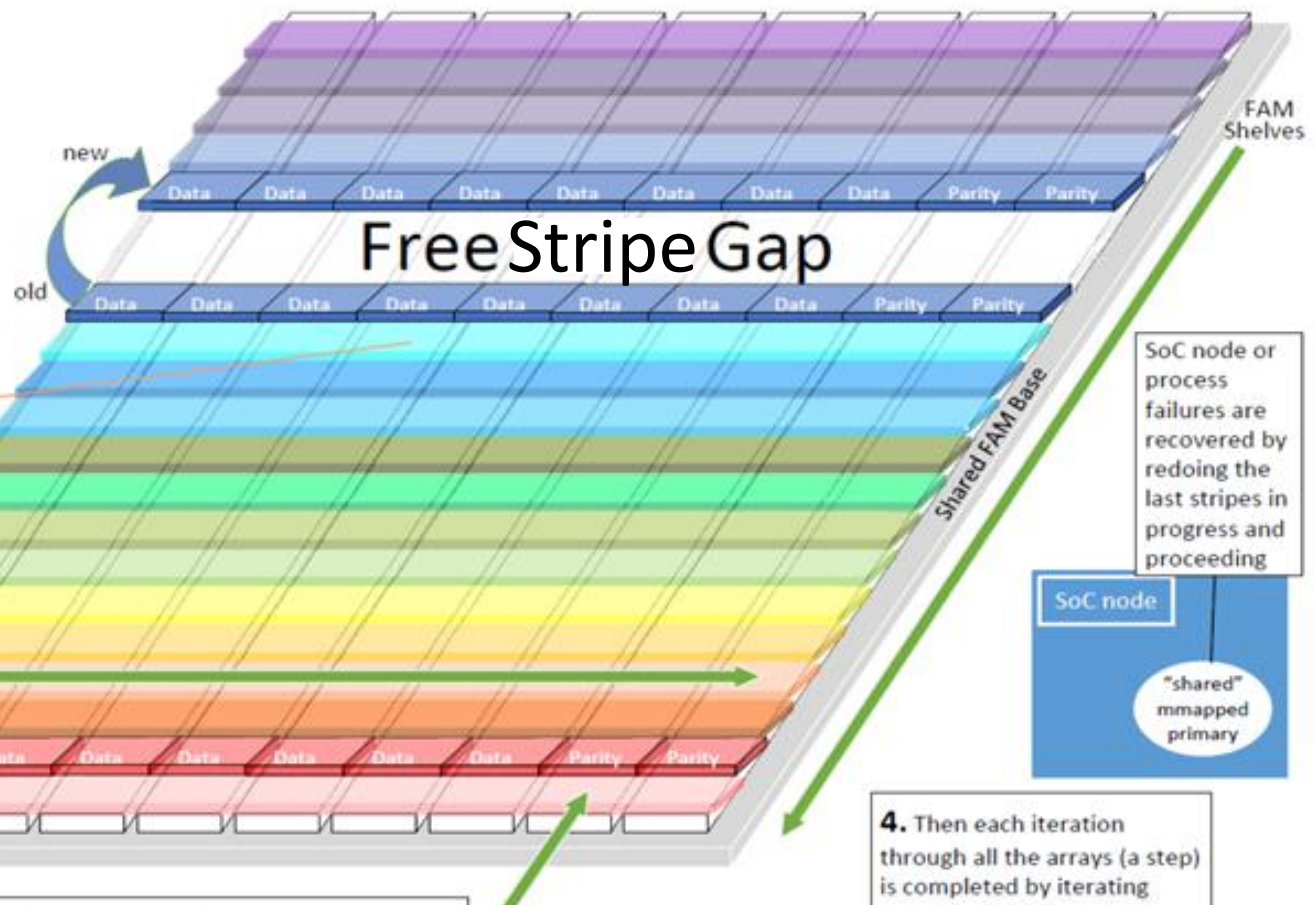
3. That leaves two to three hyperthreads (out of 4) per core for executing non-temporal (hopefully) load and store instructions

4. Then each iteration through all the arrays (a step) is completed by iterating through the stripes.

SoC node or process failures are recovered by redoing the last stripes in progress and proceeding

SoC node

"shared" mmapped primary



TxHPC presentation and code

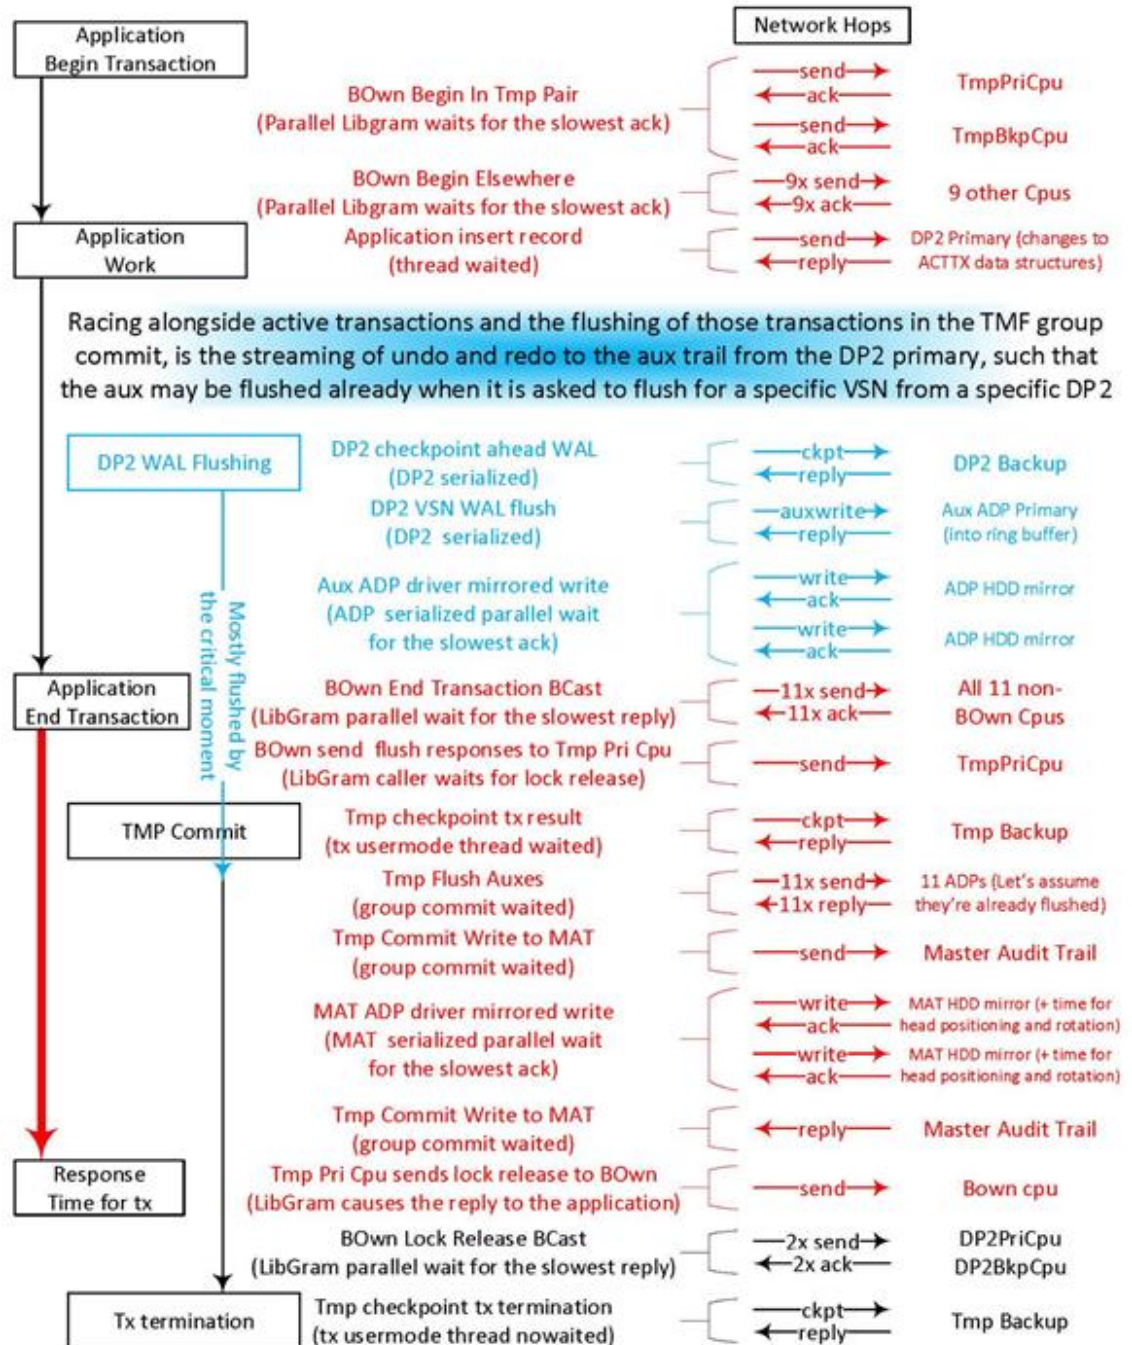
- Two-page abstract:
<http://nvmw.ucsd.edu/2017/assets/abstracts/20>
- Slides:
<http://nvmw.ucsd.edu/2017/assets/slides/20>
- Open source github:
<https://github.com/HewlettPackard/Redhead>
- TxHPC source (uses Jerasure 2.0 + GKComplete):
<https://github.com/HewlettPackard/Redhead/tree/master/include/StencilForTxHPC/TxHPC4TM>

Nonstop SQL Subtransactions

- Nonstop clustered group 3-phase commit is the slowest in the business, response time in 10s of ms. at best for standard configurations.
- They used to have 90% of the trading business, only a couple of exchanges left: this is because all flash trading completes for the front-running and arbitrage of a single trade in 15 μ s.
- In 1999 came up with a solution, presented to HPTS, but the problem wasn't pressing, then.
- Now it's an issue, so I was called in to fix the Nonstop commit code that I designed and wrote (with much help from Pat, Shel, Jimbo, Matt M., J. Carley, J. Klein, etc.)
- With SQL Subtransactions, we could get 4 orders of magnitude, with H/W work maybe another 2-3 orders in both throughput and response time (Big Ω .)

Single Record Insert on a 12 cpu Nonstop with 11 ADPs using TMF transactions

- It's completely scaled out and bullet proof.
- That translates to slow.
- There are lot of waited steps in RED.
- The part in BLUE could be very fast if we could just execute that part.
- We need a new transaction type that fits into the old transaction recovery system: SQL subtransactions.
- They need a new delivery system: a special message as a top-level transaction.
- They need to execute completely within a single disk process instance, call it a DPX.
- They need collocation to a single processor cache hierarchy to reduce response time.
- They need all resources to be confined to a single disk process.
- They need buffering/multiplexing to increase efficiency and throughput.
- They need to be ACID and the same level of high availability as TMF transactions.
- They need to be as programmable as TMF transactions, modulo the issues of closure on collocated resources.



- SQL Subtransactions has reached detailed design, 3rd revision of the spec awaiting a spot in the very busy Nonstop software development schedule (currently supporting the new Virtualized Nonstop VM on x86_64 for Gen9 hardware.)
- On to the next project that advances the state of the art of resilience.