

# **My Perspective on How SOS Got Started and Some Related Things**

**James Tomkins**

**SOS 20 Workshop  
Asheville, North Carolina  
March 22-25, 2016**


- If you noticed, the title of my talk doesn't match the agenda - The agenda title was Bill Camp's title
- Bill Camp is ill and couldn't come

# A Turning Point for Sandia

- Prior to 1987 Sandia was a follower in High Performance Computing
- In 1987 Sandia embarked on a path to leadership in High Performance Computing (HPC) through Massively Parallel Processing (MPP)


# Sandia HPC History

- **1987-- Sandia fielded first true MPP, 1024 node nCUBE 10**
  - Won Karp Challenge
  - Won inaugural Gordon-Bell Prize
  - Awarded patents for nearly every aspect of MPP software
  
- **1988-- Sandia fielded a CM2 with 16K 1 bit processors and 512 32 bit floating point units**
  - SIMD architecture
  - Later upgraded to CM-200



# Sandia HPC History

- **1990-- Sandia fielded 2 1024-node nCUBE-2's, First true MPP supercomputers**
  - Outperformed Cray vector computers @ ~1/7 the cost!
  - Sandia began research on Light Weight Kernel (LWK) operating systems
- **1993-- Sandia fielded first ~1850-node (~3900 processors) Intel Paragon**
  - Sandia (Intel) Paragon is #1 on Top 500 list (first for Sandia)
  - Wins Sandia's second Gordon-Bell Prize
  - First use of Sandia developed Light Weight Kernel (LWK) Operating System software for production computing



# Sandia HPC History

- **1997-- Sandia fielded Intel Tflops, world's first Terascale computer, 4600+ nodes (9200+ processors)**
  - Ran Sandia developed LWK System Software
  - Number 1 on the “Top 500” list for 7 consecutive lists from June, 1997 through June, 2000, a record still unmatched!
- **1997-- Sandia began development of world's first Linux “Super-cluster”, Cplant™**
  - Sandia integrated DEC/Compaq HW with Myrinet network
  - Sandia developed all run-time, file system and messaging software
  - World's first terascale cluster-- Cplant grew to become the world's first terascale cluster and achieved ~1 TF on Linpack in 2003.



# Some History

- Sandia forms a partnership to acquire an Intel Paragon in 1992
  - No more PVP machines
  - Funding - AFWL (DARPA) and internal Sandia funding
- Bill Camp takes a temporary position at Cray in early 1994
- Sandia (Ed Barsis with Al Narath's support) makes a deal with DOE in the fall of 1994 to get the first one of the ASCI Supercomputers. And in September of 1995 Sandia signed a contract with Intel to build ASCI Red.
- Sandia Management changes in the fall of 1995
  - Al Narath is promoted to Lockheed Martin Sector VP
  - Ed Barsis retires from Sandia at the end of 1995

# More History

- Bill Camp returns to Sandia from Cray in the fall of 1995 and in early 1996 he is promoted to Director of Advanced Computing at Sandia.
- SGI buys Cray in early 1996 - No T3F
- Intel decides to get out of supercomputing in early 1996 - Gil Weigand is furious
- Bill Camp, Rolf Riesen, and I travel to Switzerland in April of 1996 (Meet with Ralf Gruber and others at EPFL and ETH)
  - The Swiss show a strong interest in building clusters
  - The Swiss also indicate a desire to have a technical relationship with Sandia
- Bill and I meet with Ken Kliwer of ORNL at SC96 to discuss starting a workshop with Sandia, ORNL, and the Swiss as the organizers
  - Sandia and ORNL were both managed by Lockheed Martin in 1996



# Cluster Computing Driver for Sandia

- Sandia was being cut out of large ASCI Machines
- Sandia needed a way to continue at the forefront of parallel computing
- Cplant was part of Sandia's answer
  - Commodity Components
  - Custom Software

# Results

- Sandia began work on Cplant in early 1997
- Sandia organized first SOS Workshop in Santa Fe, NM in late summer of 1997
- Other Sandia Initiatives - Changes in Sandia and DOE Management (1999-2000)
  - Red Skunk - ASCI Red upgrade
  - Red Storm (Cray XT)

# Some Comments

- Building a “supercomputer” out of commodity parts was hard
  - Components did not always work as advertised
  - Building and supporting a full set of system software was more difficult than expected
- Clusters had potential - Vendors got involved
  - Capacity Computing
  - Google, etc.
- Clusters were not equivalent to a true MPP

# **A Final Comment**

- Politics can be a really important factor that can not be ignored



Bobcat Family