

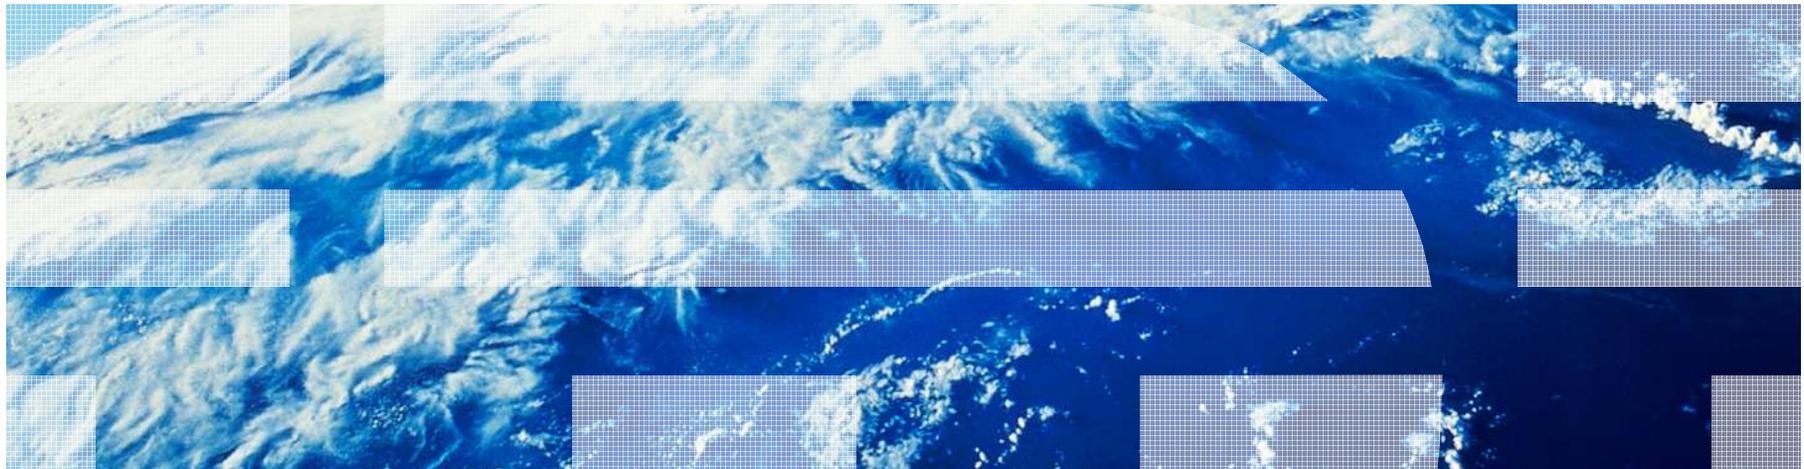


Robert W. Wisniewski  
Chief Software Architect  
Blue Gene Supercomputer Research

November 15, 2011



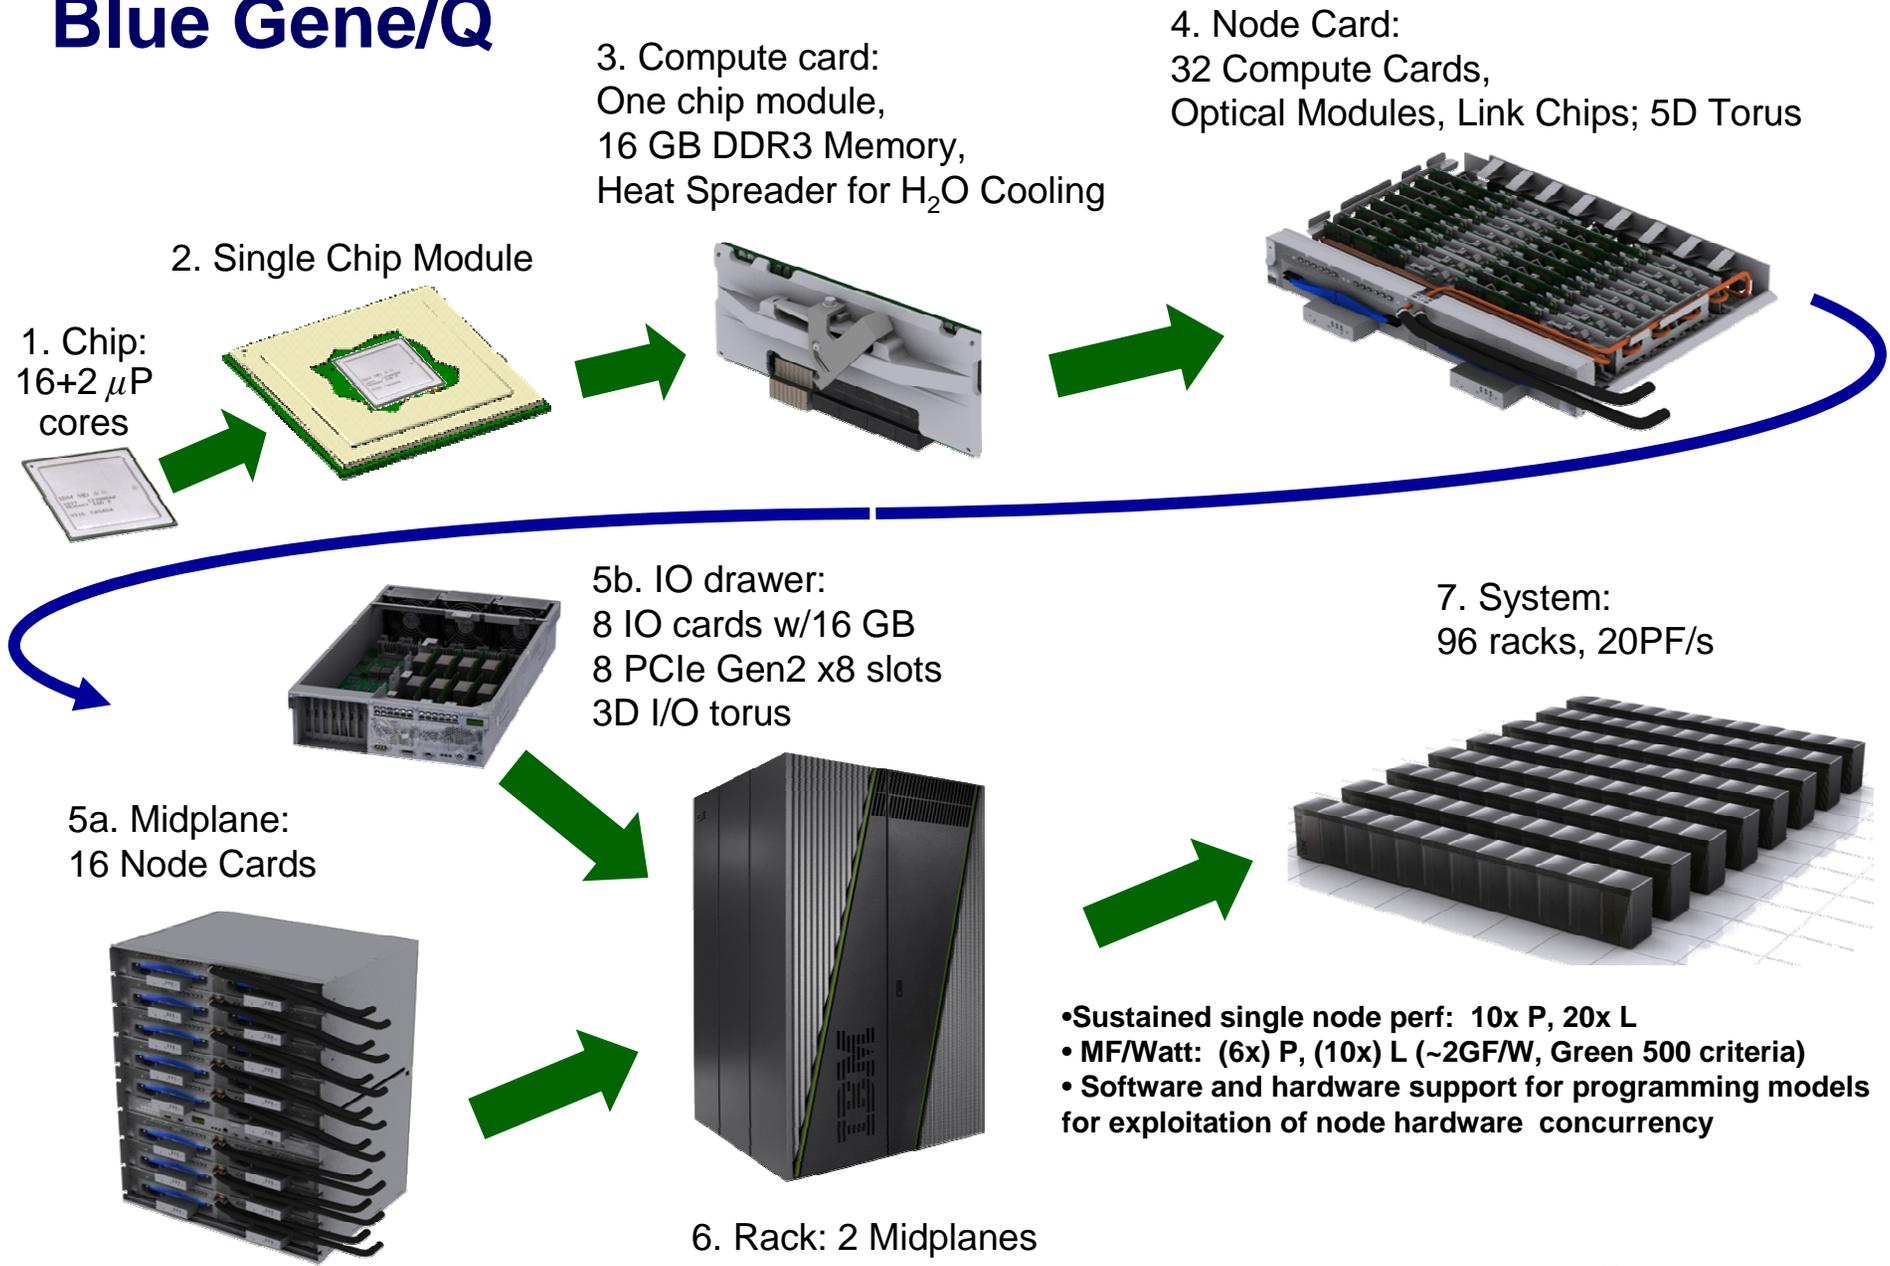
# Codesign on MPICH



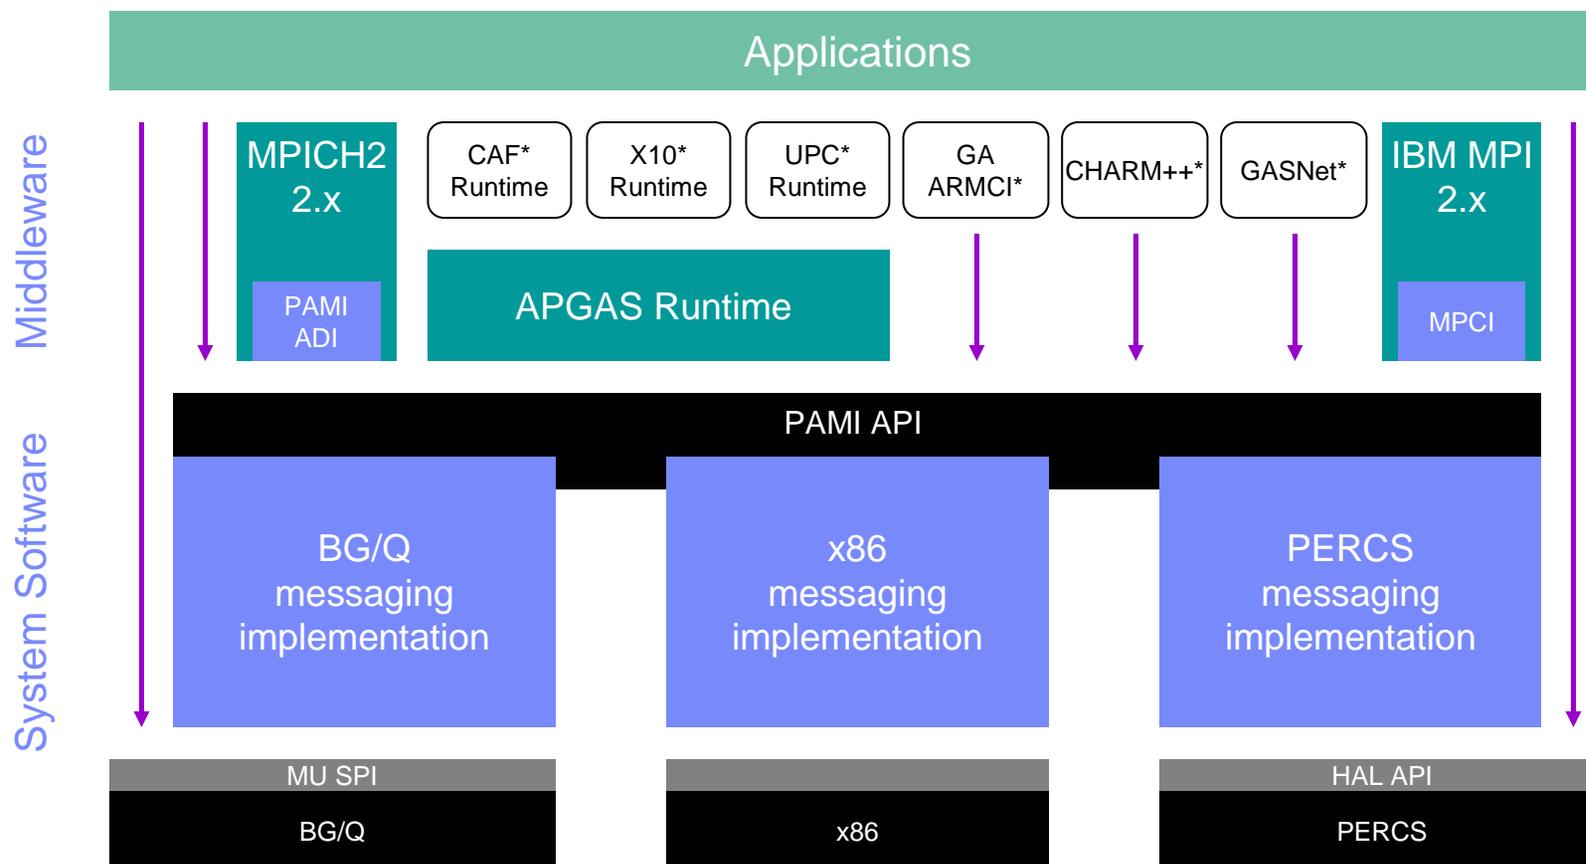
## Collaboration with MPICH Team for BG/Q

- Began nearly 4 years ago for BG/Q
- Observation of where Q architecture was going to be 16 cores x 4 SMT
- Challenge both with intranode and internode scalability
  - Intranode
    - Each node is 16 cores 4-way SMT = 64 hardware threads
    - Supporting up to 64 MPI tasks per node
  - Internode
    - Sequoia 98304 nodes x 64/node = 6291456 (6.3M) MPI tasks
- Partnered with MPICH team across messaging stack
  - MPI, PAMI, and SPI

# Blue Gene/Q

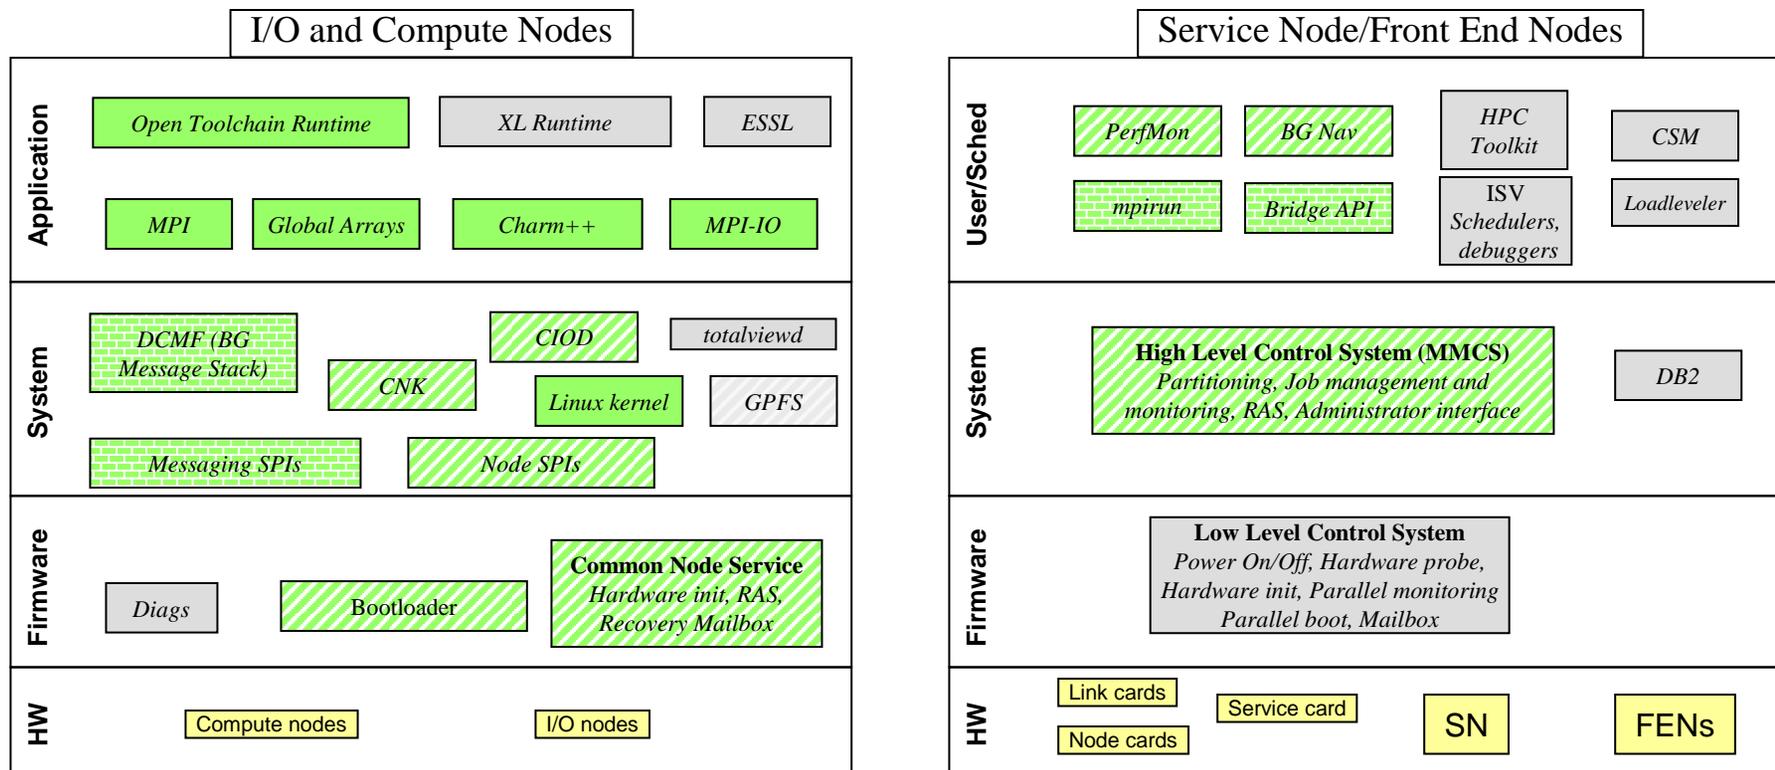


# Supporting Different Shared Memory/Threading Models on PAMI Can Have Familiar Programming Models for Exascale



\*all runtimes are not supported on all platforms

# Blue Gene Software Stack Openness



- New open source reference implementation licensed under CPL.
- New open source community under CPL license. Active IBM participation.
- Existing open source communities under various licenses. BG code will be contributed and/or new sub-community started..
- Closed. No source provided. Not buildable.
- Closed. Buildable source available

# Outcome

- Will release Blue Gene/Q software early next year with MPICH
- Messaging rates and latency meeting goals
- Expand to support across high-end IBM HPC
- The expected integration of Platform Computing technologies will augment not replace capability, for example, we will continue to pursue MPICH as a high-end solution across IBM HPC