The Roots of Beowulf

James R. Fischer
NASA Goddard Space Flight Center
Code 606
Greenbelt, MD 20771 USA
1-301-286-3465
james.r.fischer@nasa.gov

ABSTRACT

The first Beowulf Linux commodity cluster was constructed at NASA's Goddard Space Flight Center in 1994 and its origins are a part of the folklore of high-end computing. In fact, the conditions within Goddard that brought the idea into being were shaped by rich historical roots, strategic pressures brought on by the ramp up of the Federal High-Performance Computing and Communications Program, growth of the open software movement, microprocessor performance trends, and the vision of key technologists. This multifaceted story is told here for the first time from the point of view of NASA project management.

Categories and Subject Descriptors

C.1.4 [Processor Architectures]: Parallel Architectures --- distributed architectures; C.5.0 [Computer System Implementation]: General; K.2 [Computing Milieux]: History of Computing

General Terms

Management

Keywords

Beowulf, cluster, HPCC, NASA, Goddard, Sterling, Becker, Dorband, Nickolls, Massively Parallel Processor, MPP, MasPar, SIMD, Blitzen

1. INTRODUCTION

Looking back to the origins of the Beowulf cluster computing movement in 1993, it is well known that the driving force was NASA's stated need for a gigaflops workstation costing less than $50,000. That is true, but the creative conversations that brought the necessary ideas together were precipitated by a more basic need—to share software.

2. THE PRE-BEOWULF COMPUTING WORLD

A flashback to the pre-Beowulf computing world paints a picture of limitations. The perspective is NASA centric, Goddard Space Flight Center specifically, but the experience was universal. It is only 20 years ago, but the impediments facing those who needed high-end computing are somewhat incomprehensible today if you were not there and may be best forgotten if you were.

20 Years of Beowulf: Workshop in Honor of Thomas Sterling's 65th Birthday, October 13–14, 2014, Annapolis, Maryland, USA.

2.1 Proprietary Stove Piped Systems

Every system that we could buy ran proprietary system software on proprietary hardware.

Competing vendors' operating environments were, in many cases, extremely incompatible—changing to a different vendor could be traumatic.

A facility's only recourse for software enhancement and problem resolution was through the original vendor.

2.2 Poor Price/Performance

In 1990, "a high-end workstation can be purchased for the equivalent of a full-time employee." [1]

In 1992, some facilities used clusters of workstations to offload supercomputers and harvest wasted cycles; at Goddard we salivated at this idea but had few high-end workstations to use.

The very high and rising cost of each new generation of supercomputer development forced vendors to pass along those costs to customers. The vendors could inflate their prices because they were only competing with other vendors doing the same thing.

2.3 Numerous Performance Choke Points

In 1991, the Intel Touchstone Delta at Caltech was the top machine in the world, but compilation had to be done on Sun workstations using proprietary system software that only ran on Suns.

In 1993, Connection Machine Fortran compile and link performance averaged about 10 lines per second on the host; performance was similar for the MasPar used at Goddard. All development had to be done on the host machine; vendors were not really solving this problem (maybe they could sell you two host machines).

2.4 Instability

In 1993, operational metrics recorded by NASA Ames Research Center for their Intel Paragon reported "Reboots Weekly Average" (typically 15–30) and "Mean Time to Incidents" (typically 4–10 hours). Each reboot forced all running jobs to be restarted, and the reboot for some systems might take 30 minutes. Since the bigger MIMD machines were usually one-off's, the OS developers had to take the entire system away from users into