

# BIOVULF

HIGH PERFORMANCE COMPUTING AT THE NIH

20<sup>th</sup> Anniversary  
1999-2019

# SC97 Tutorials, San Jose



Sunday, November 16  
Full Day 8:30am-5pm

## **How to build a Beowulf: Assembling, Programming, and Using a Clustered PC Do-it-yourself Supercomputer**

**Thomas Sterling, California Institute of Technology**

It has recently become possible to assemble a collection of commodity mass market hardware components and freely available software packages in a day and be executing real world applications by dinner time to achieve a sustained performance at greater than 1 Gflops at a total cost of around \$50,000. Furthermore, on almost a daily basis, these numbers are improving. This full-day tutorial will cover all aspects of system assembly, integration, software installation, programming, application development, system management, and benchmarking. Demonstrations with actual hardware and software components will be conducted throughout the tutorial. Participants will be encouraged to closely examine and manipulate elements of a Beowulf at various stages of integration with strong Q&A interaction between presenters and attendees. The presenters will include David Bailey (applications benchmarking), Don Becker (networking), Jack Dongarra (applications benchmarking), Al Geist (PVM), Ewing (Rusty) Lusk (MPI), John Salmon (applications), and Thomas Sterling (system structure).


# Biowulf 1999



- Compute cores: 80
- GPUs: huh?
- Fileservers: 2
- Storage: Gigabytes
- Networking: 100 Mb/s
- Applications: 2
- Staff: 2
- Publications: 0

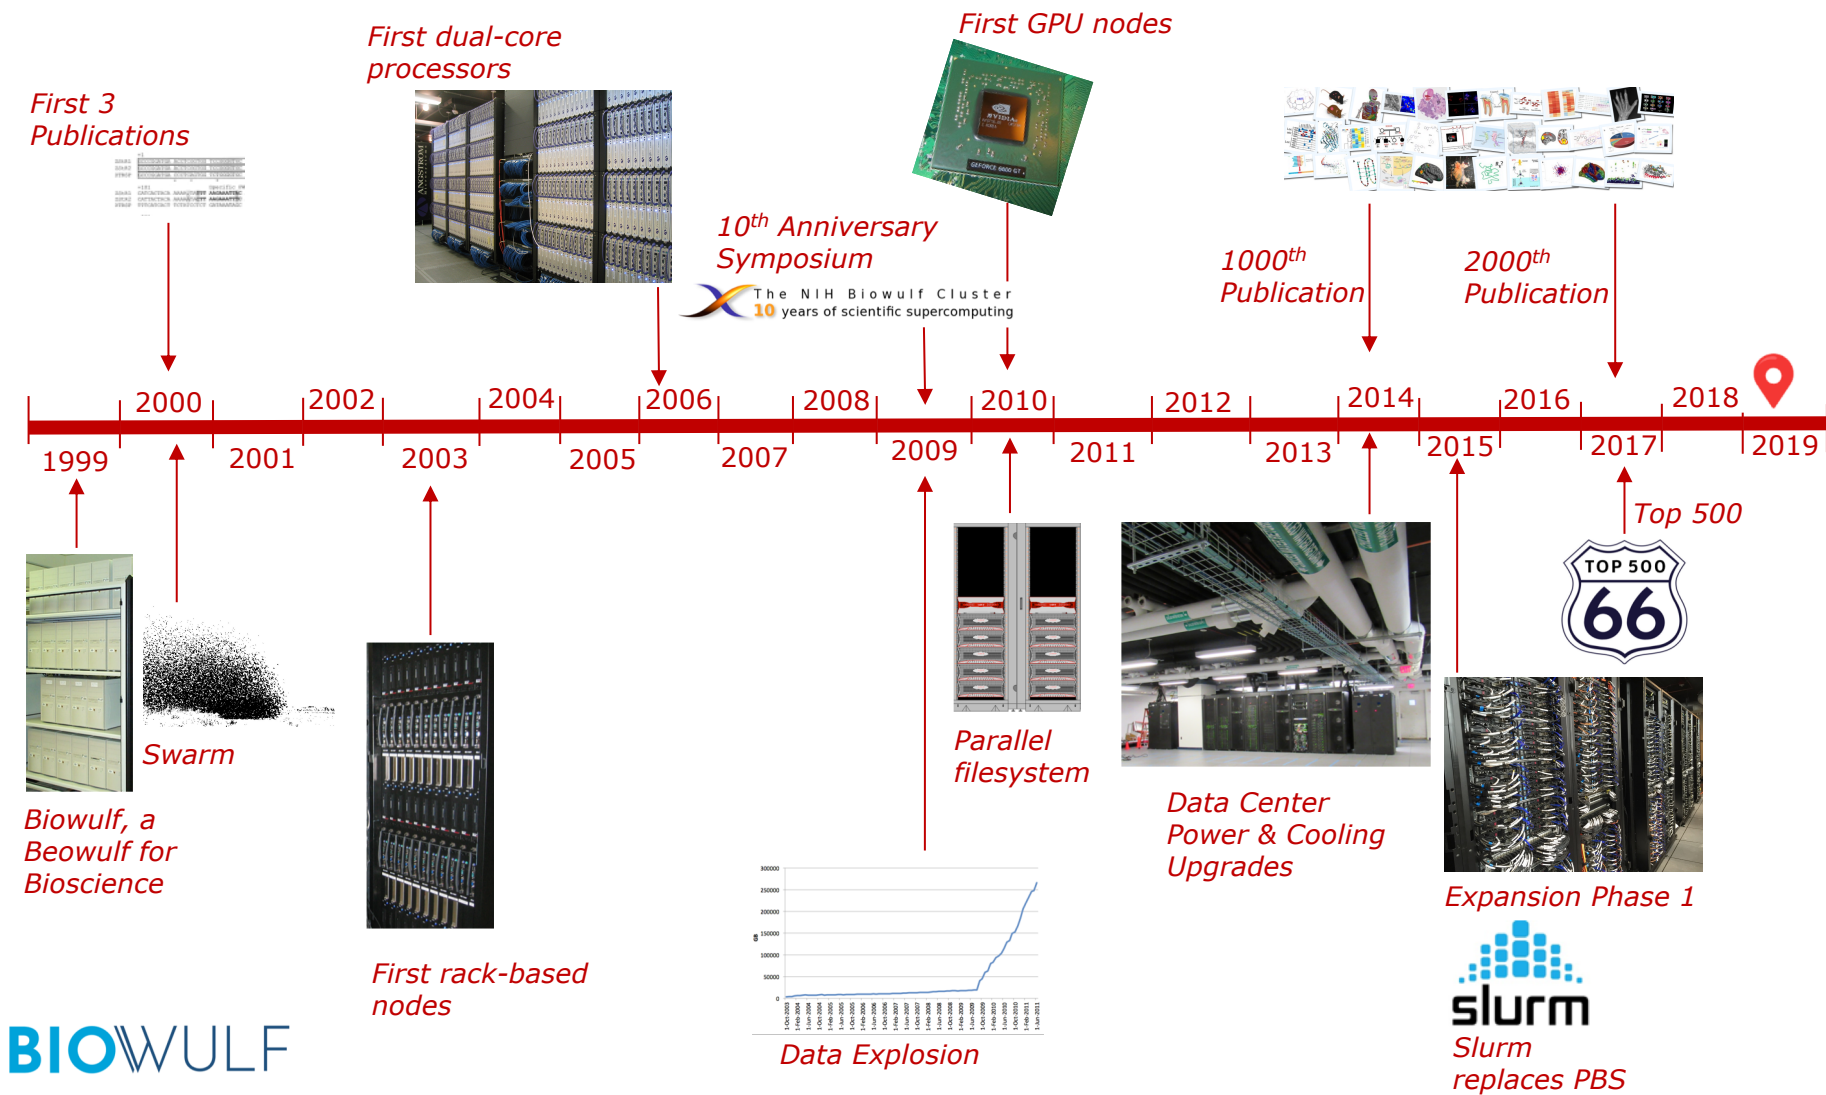
# Biowulf 2019

- Compute cores: 95,000
- GPUs: 560
- Fileservers: 82
- Storage: Petabytes
- Networking: 100 Gb/s
- Applications: 661
- Staff: 18
- Publications: 2486






# Getting from There to Here...




# Twenty Years Growth

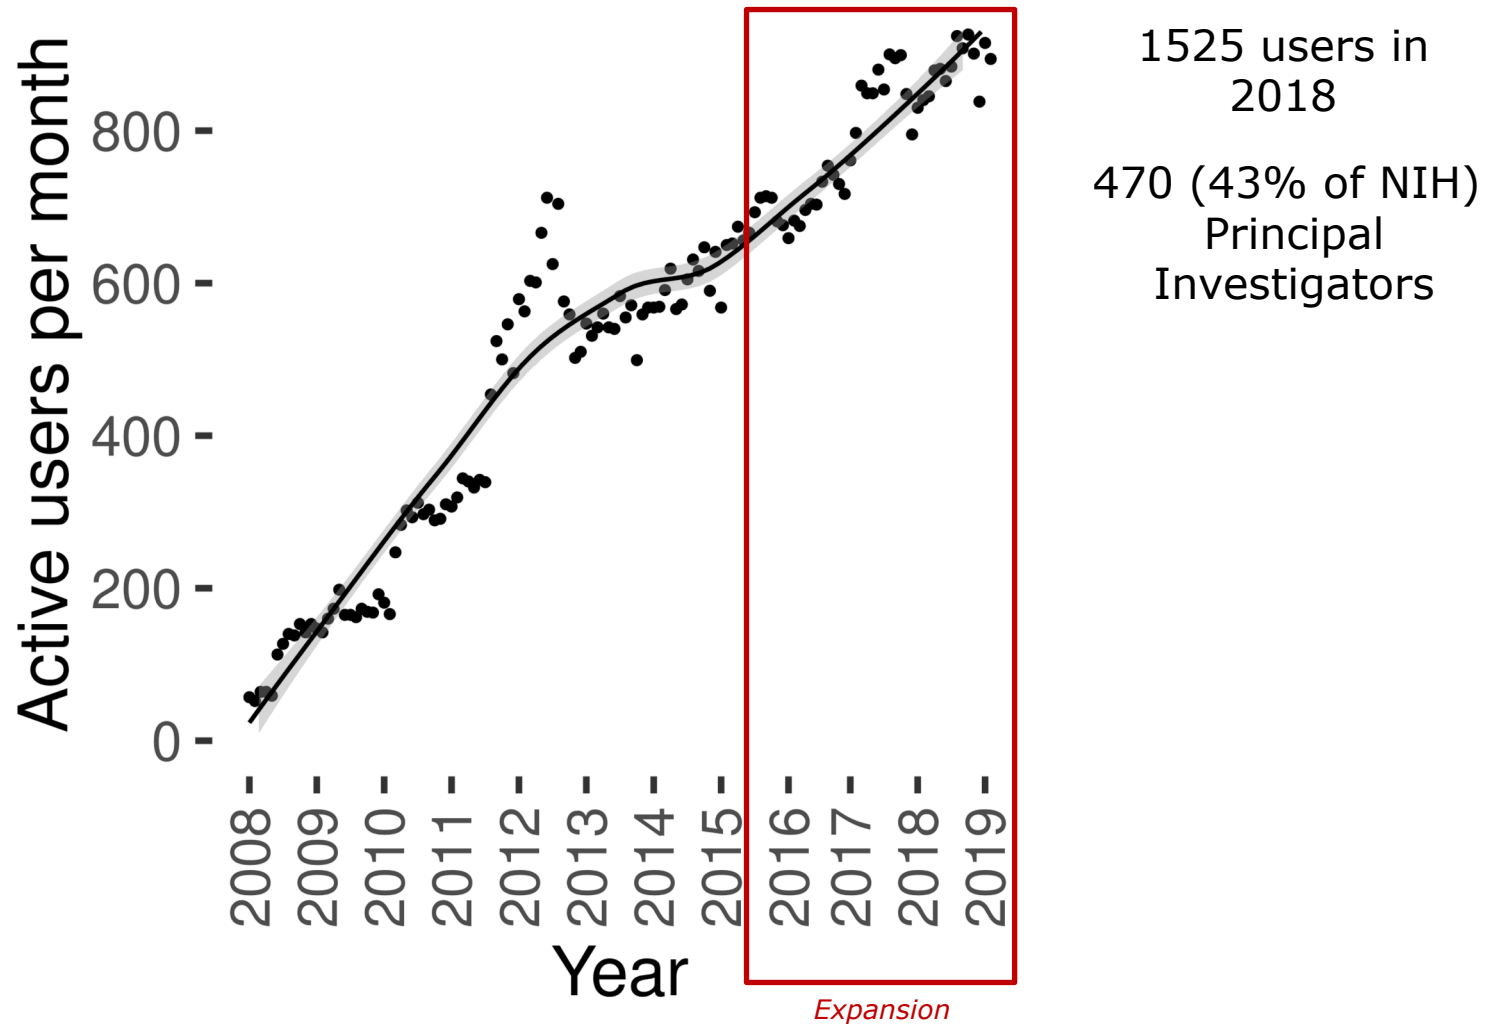
## Compute Cores




## Publications Using Biowulf



# Active Users




# CPUh Used





# Data Storage

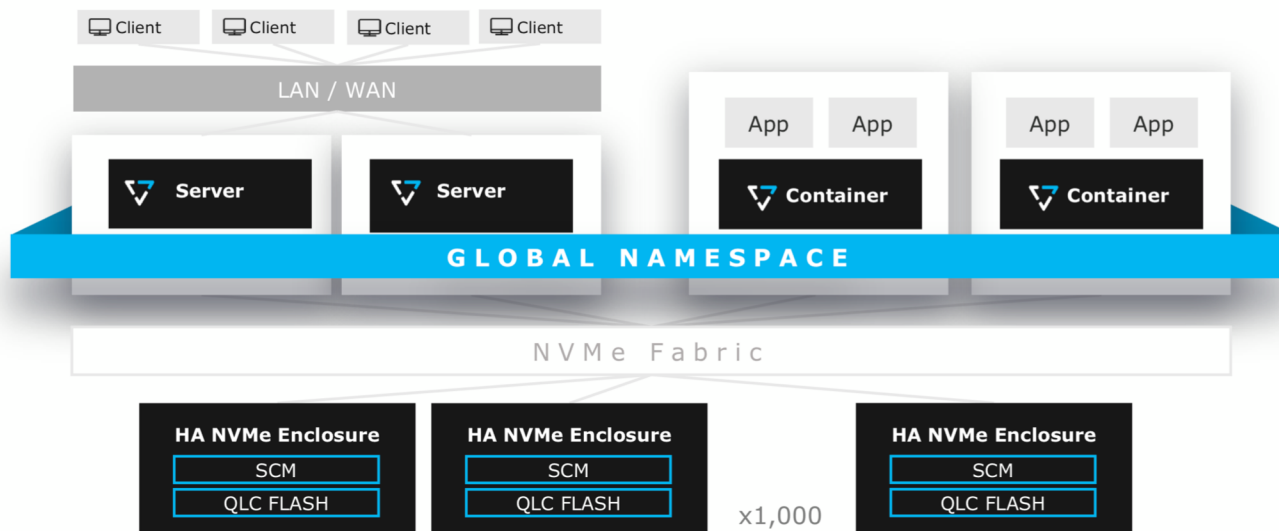


## BIO\WULF

[illegible]

# Phase 5: Coming Summer 2019

- 6,912 “Skylake” compute cores (384 GB/node)
- 224 V100 GPUs (32 GB, NVLINK)
- 100-200 Gb/s Infiniband networking
- Solid State network-based shared storage



# Biowulf Staff 1999-2019

---





# BIOVULF

HIGH PERFORMANCE COMPUTING AT THE NIH