



# Linux HPC Software Stack

**Makia Minich**

Clustre Monkey, HPC Software Stack

Lustre Group

April 2008



# Project Goals

- Develop integrated software stack for Linux-based HPC solutions based on Sun HPC hardware
- Provide a complete set of tools and well-defined processes for configuring and provisioning an HPC cluster
- Design the stack to be scalable
- Include tools for verification, management and administration, and monitoring

# About the Project

- “Open” and community-driven
  - > “Open” project - not a corporate product
  - > Using as much open source as possible
  - > Disseminating information (community style)
  - > Looking for community input on how to handle things better
- Built on existing Linux distributions
  - > HPC options are modular on top of the OS
- Includes a verification suite
  - > Does our stack install and behave correctly?
  - > Will 3<sup>rd</sup> party packages work with our stack?

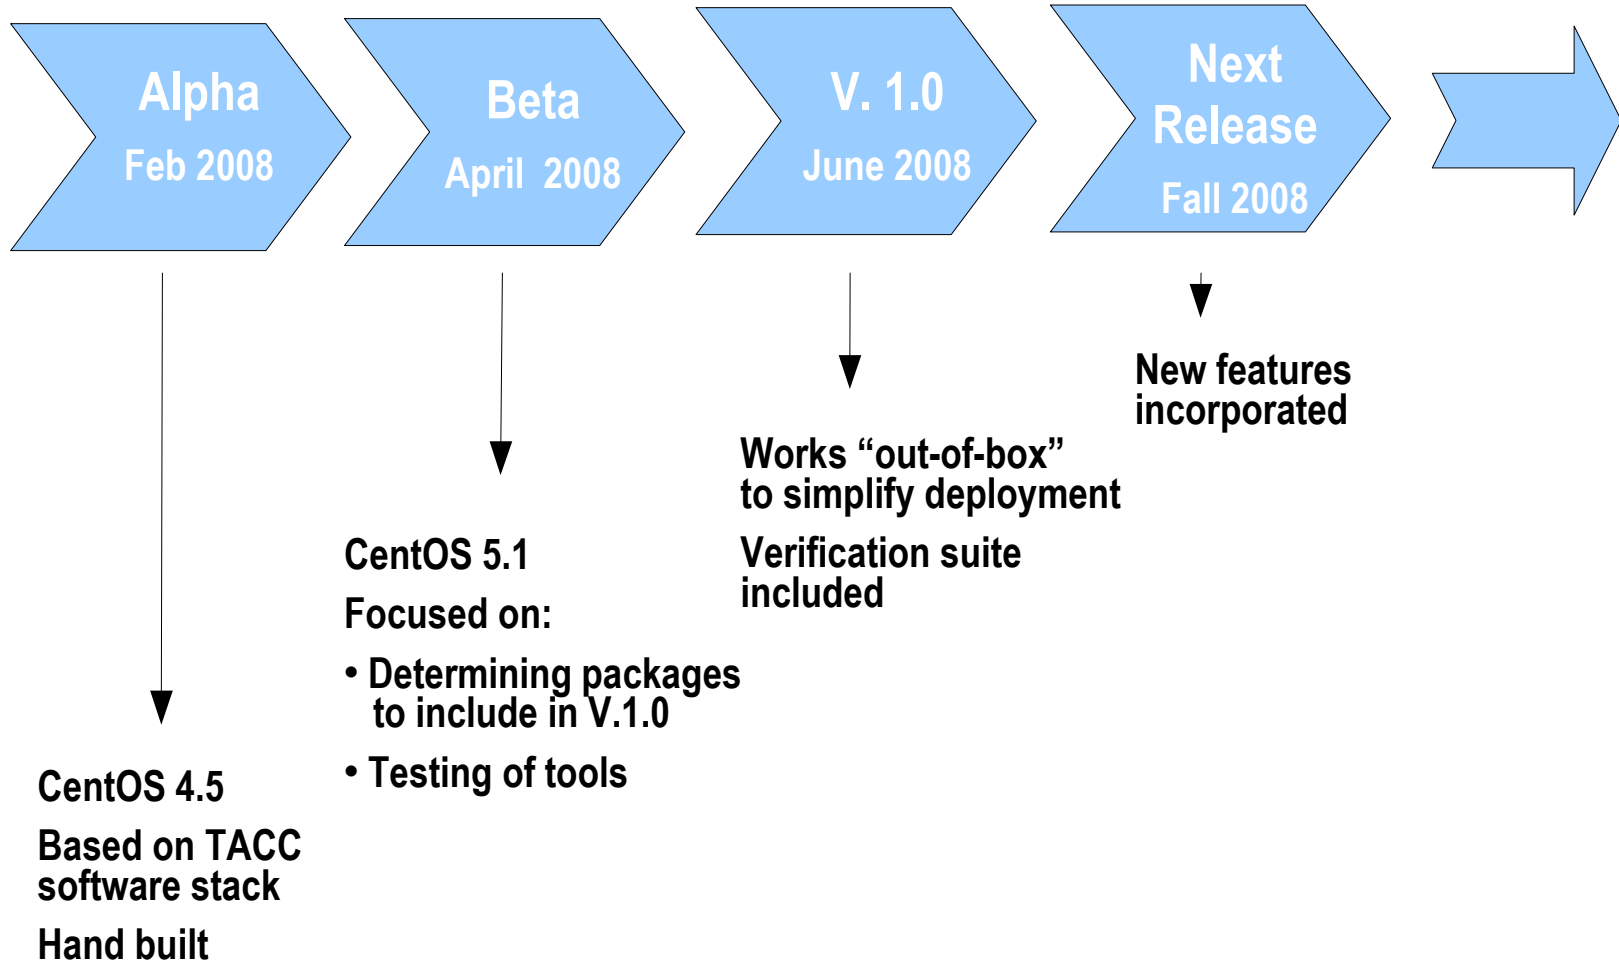
# About the Team

- Currently seven team members
- Areas of expertise include:
  - > Sun architecture and system software
  - > Lustre and Linux
  - > System integration
  - > Technical communications
  - > IB specialist, quality engineer (to be hired)
  - > Interns – working on testing and other side tasks
- Distributed across U.S., Canada, and Japan

# Project History

- Project announced at SC07 in November 2007
- Team rapidly formed in January 2008
- Project officially launched mid-February 2008
- Beta wrapping up now
- On track for first release in June 2008

# Linux HPC SW Stack Roadmap



# About the Stack

- Base distro choice based on community input
- Verification suite helps ensure inter-operable 3<sup>rd</sup> party software
  - License may be required to include 3<sup>rd</sup> party software in stack
  - Verification suite will be enhanced with new tests as issues are resolved to ensure old bugs stay dead and new ones are flushed out
- Product delivered as download or on DVD
  - Contains the component framework needed to start turning a bare-metal system into a running HPC cluster



*SUN*<sup>®</sup>  
microsystems



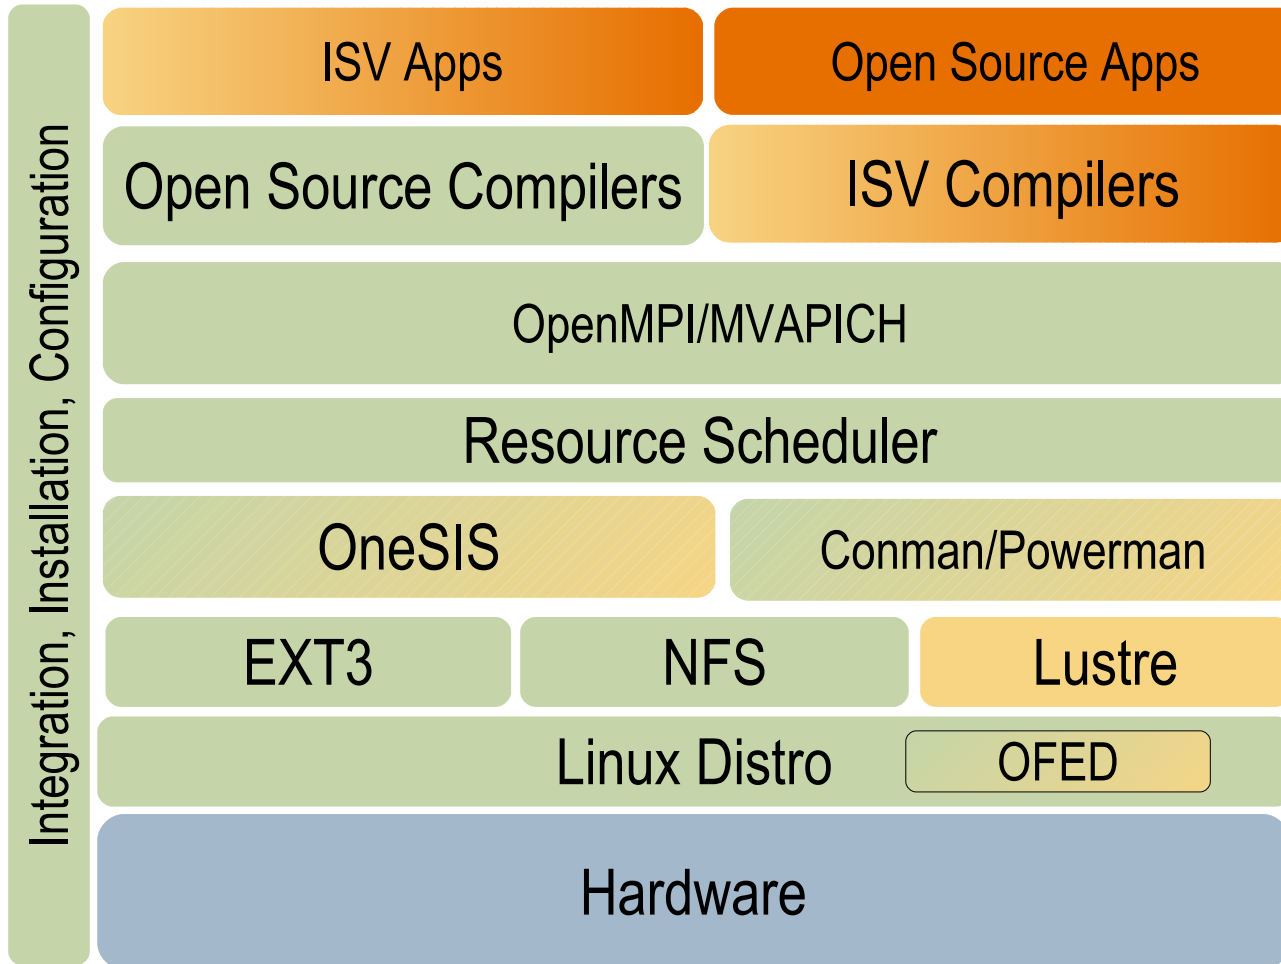
Welcome to the Sun HPC Software Stack Installer.

- To install, press the <ENTER> key.
- To perform a memory test, type memtest <ENTER>

boot: \_



# Linux HPC Software Stack



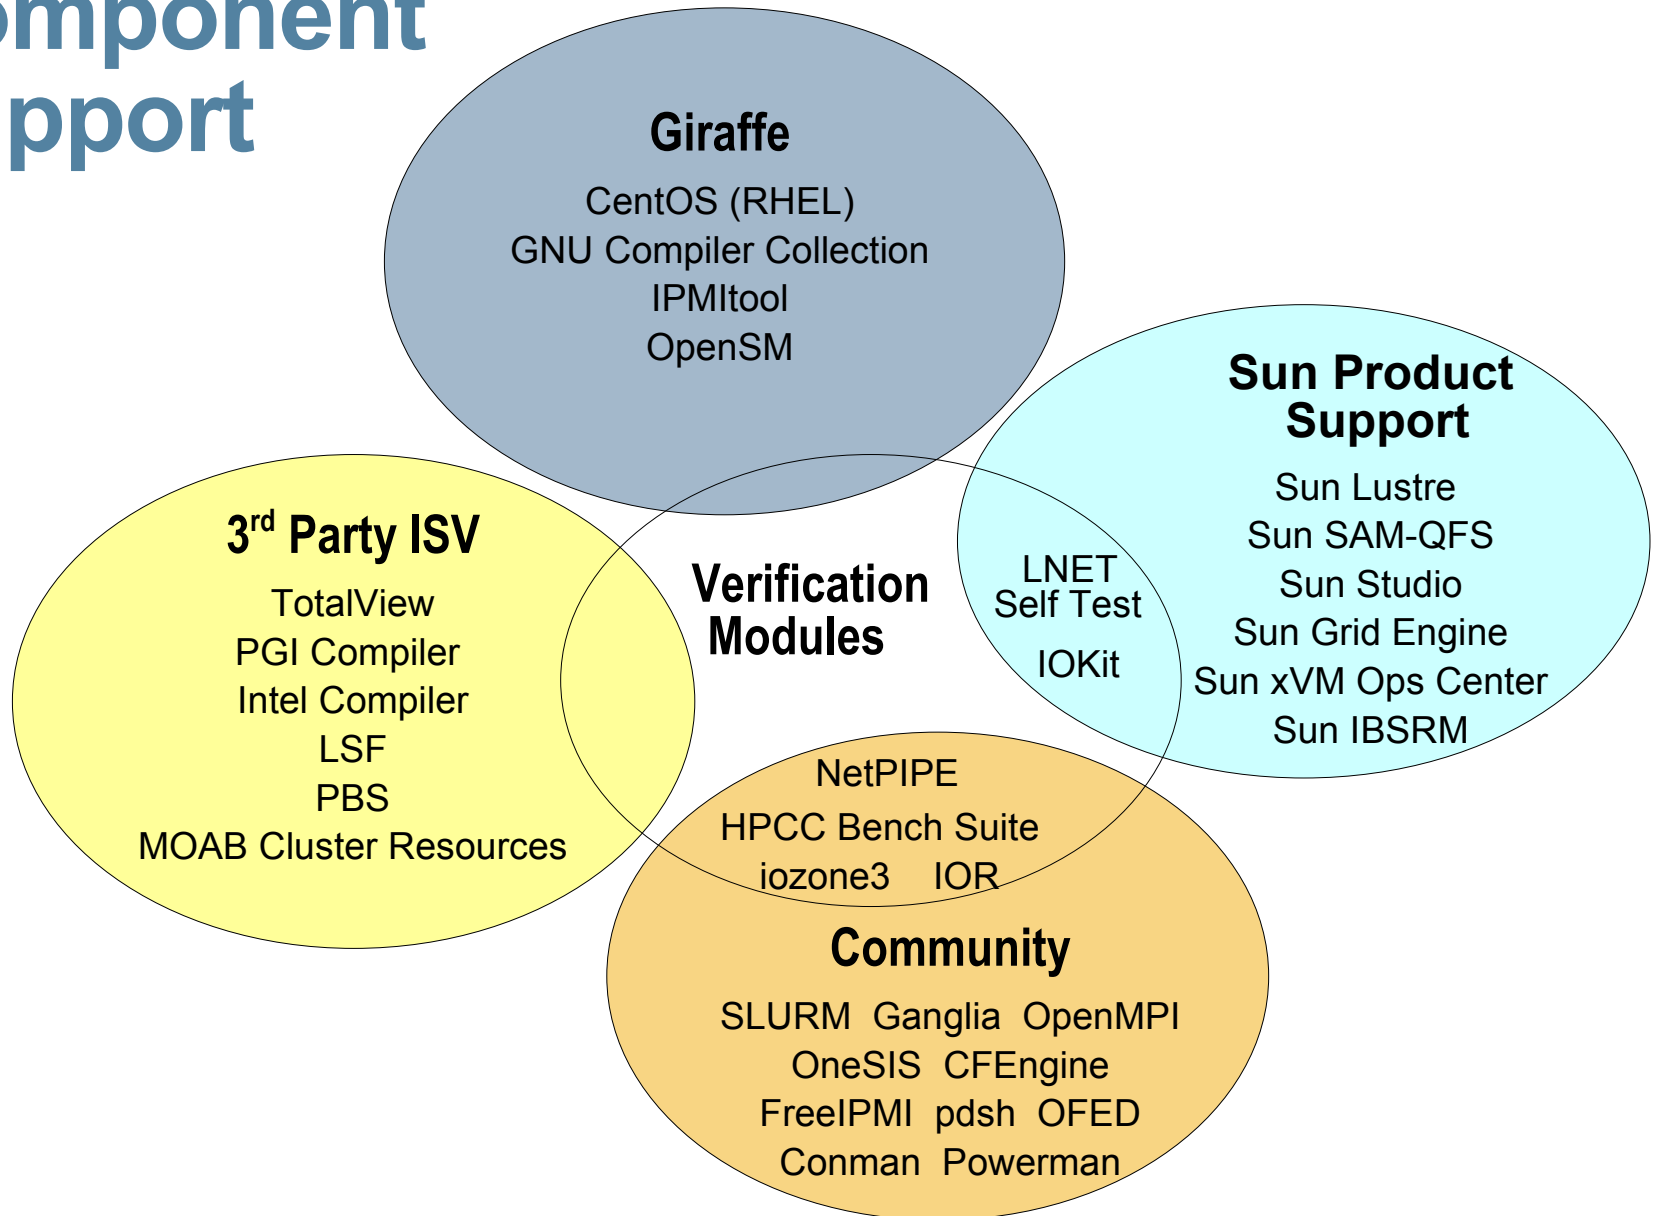
# Linux HPC SW Stack Components

Provisioning	OneSIS									
Schedulers	<u>Slurm</u>	<b><i>MOAB</i></b>	<u>SGE</u>	<u>PBS</u>	<u>LSF</u>					
OS/ User space/ Apps/ Verification Modules	lozone3	IOR	NetPipe	HPCC Benchmark	IOkit	LNET Self Test				
OS/ User space/ Compilers & Debuggers	<b><i><u>Intel Compilers</u></i></b>	<b><i>OpenMPI</i></b>	<b><i>PGI</i></b>	<b><i>Totalview</i></b>	Gnu Debuggers	<u>SunStudio</u>	<u>Modules</u>	<u>Heartbeat</u>		
OS/ User space/ Libraries	OFED Libraries	Open MPI	OFED 1.2.x	<u>LaPACK</u> , <u>Atlas</u>	<u>AMD's ACML</u>					
OS/ Kernel space	CentOS	Lustre	OSS/MDS	<u>Sam-QFS</u>						
Adapters	Ethhernet	Mellanox IB	SATA	Fibre Channel	Graphics/ VIZ	GPGPU				
(Server) HW platform	SW RAID	Constellation	Galaxy/ blade/etc. technologies	Server & switch (Voltaire / Magnum)	Other server switches	Drivers including OFED	Firmware upgrades	Component chips	HW counter	HW management (network gear)
Node	ILOM	IPMI-based	Service processors	Power on/off	BIOS integration					

\* Bold and Italics means licensed software, validated but not included. Underlined item is an option a customer may wish to add.

	HPC SW Stack components
	Beta
	V1.0 (under consideration, includes Beta components)

# Component Support



# Communication Paths

- Email List
  - > [linux\\_hpc\\_swstack@lists.lustre.org](mailto:linux_hpc_swstack@lists.lustre.org)
- Blog
  - > <http://blogs.sun.com/giraffe>



**Makia Minich**  
makia@sun.com

