



Bootable Cluster CD Supercomputing 2011

Ivan Babic

Earlham College

`ibabic09 at cs.earlham.edu`

Tom Murphy

Contra Costa College

`tmurphy at contracosta.edu`

Andrew Fitz Gibbon

Shodor Foundation

`fitz at cs.earlham.edu`

Charlie Peck

Earlham College

`charliep at cs.earlham.edu`

Aaron Weeden

Earlham College

`amweeden06 at cs.earlham.edu`

Mobeen Ludin

Earlham College

`mmludin08 at cs.earlham.edu`

Skylar Thompson

University of Washington

`skylar at cs.earlham.edu`

November 14, 2011



Sponsors

Sponsors

Preliminary work

What is the BC CD?

Short history

What is the problem?

How is BC CD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

- Intel Corporation
- Supercomputing Conference
- Earlham College
- Shodor Foundation
- Contra Costa College



Preliminary work

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

1. Insert the USB stick.
2. Open a terminal (button by Firefox at the top).
3. Type in:

```
wget -O - \
http://cluster.earlham.edu/bccd-ng/sc11/make-bccd| \
bash
```



What is the BCCD?

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

- A software tool for domain educators for teaching CSE
- A stable platform for the writing of curricula for CSE
- A live CD with pre-configured clustering software
- One part of a larger CSE community

Shodor Foundation

National Computational Sciences Institute (NCSI)

CSE Reference Desk (CSERD)

- “Teaching Parallelism Made Easy”

Just keep pushing “Enter”



Short history

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

- First versions developed by Paul Gray and Students at University of Northern Iowa
- Current version (BCCDv3) collaboratively developed by a number of educational institutions
- BCCDv3 much easier to maintain and able to respond to new needs
- BCCDv3 has the same user experience as BCCDv2, so curricula written for v2 will still work with v3



What is the problem?

Sponsors
Preliminary work
What is the BCCD?
Short history
What is the problem?
How is BCCD a solution?
Live Demo!
Live Demo! (cont'd)
Boot
GalaxSee - Single node
GalaxSee - Single node (cont'd)
GalaxSee - Multi-node
GalaxSee - Multi-node (cont'd)
Miscellaneous information
References
Questions (and answers)

- STEM education is more important than ever [1]
- Computers have become a large part of any science curriculum
- BUT ... it's hard to find suitable resources for teaching computational science and parallel programming at scale
- Many schools do not have the resources or expertise to provide the hardware or software
- Those that do tend to provide the resources primarily for research and not for education
- Even with a dedicated educational resource, it tends not to be “hands-on”, and is difficult to see how it works.



How is BCCD a solution?

- Sponsors
- Preliminary work
- What is the BCCD?
- Short history
- What is the problem?
- How is BCCD a solution?**
- Live Demo!
- Live Demo! (cont'd)
- Boot
- GalaxSee - Single node
- GalaxSee - Single node (cont'd)
- GalaxSee - Multi-node
- GalaxSee - Multi-node (cont'd)
- Miscellaneous information
- References
- Questions (and answers)

- Almost every school will have a computer lab. BCCD is designed to be used on existing systems non-disruptively.
- Many laptops are now dual-core. Virtualization software can be used to setup a cluster right in front of you.
- Regardless of how it's used, BCCD provides an environment that facilitates education.



Live Demo!

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

- Double click on Oracle VM VirtualBox.
- Go to New, and click Next
- Type in “BCCD” as the machine name.
- Select “Linux”, then make sure “Debian” is selected
- Enter 1024 MB for the memory.
- Unselect Start-up Disk, and click Continue when prompted.
- Click Create.
- Click Storage.



Live Demo! (cont'd)

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

- Click Empty by the CD symbol, click the CD drop-down on the right, and select Choose.
- Click on the USB stick symbol on the left, and go into SC11Materials.
- Select bccd.iso.
- Go to System, and select Processor, and set the number of processors to 2.
- Go to Network, set NAT to Internal, and set the name to “bccd”.
- Click Ok.



Boot

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

- Make sure BCCD is selected, and click Start.
- Just keep pressing Enter.

GalaxSee - Single node

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

- In the black terminal window, run `top`

- In the blue terminal window -

```
$ cd $HOME/GalaxSee
```

```
$ hostname > machines
```

```
$ make
```

GalaxSee - Single node (cont'd)

- Sponsors
- Preliminary work
- What is the BCCD?
- Short history
- What is the problem?
- How is BCCD a solution?
- Live Demo!
- Live Demo! (cont'd)
- Boot
- GalaxSee - Single node
- GalaxSee - Single node (cont'd)**
- GalaxSee - Multi-node
- GalaxSee - Multi-node (cont'd)
- Miscellaneous information
- References
- Questions (and answers)

■ Still in the blue terminal window -

```
$ time mpirun -np 1 ./GalaxSee 500 400 5000
```

```
$ time mpirun -np 2 ./GalaxSee 500 400 5000
```

Notice the number of GalaxSee processes running in the top window

You might not see speedup in VirtualBox, because of virtualization overhead.



GalaxSee - Multi-node

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

- Back in VirtualBox, right-click on the BCCD system, and select clone.
- Set the name to BCCD2, click Reinitialize, Next, and then Clone.
- Click Start for BCCD2, and follow the same boot process as the original.

GalaxSee - Multi-node (cont'd)

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)

```
■ $ module purge && module load modules mpich2
```

```
■ $ make clean && make
```

```
■ $ bccd-snarfhosts
```

```
■ $ bccd-synkdir --ni $HOME/GalaxSee \  
    $HOME/machines
```

```
■ $ time mpirun -np 4 \  
    /tmp/$(hostname -s)-$(whoami)/GalaxSee \  
    500 400 5000
```

■ Notice that not all your processes are shown in the local top window.

■ Once again, speedup is limited by virtualization.



Miscellaneous information

Sponsors
Preliminary work
What is the BCCD?
Short history
What is the problem?
How is BCCD a solution?
Live Demo!
Live Demo! (cont'd)
Boot
GalaxSee - Single node
GalaxSee - Single node (cont'd)
GalaxSee - Multi-node
GalaxSee - Multi-node (cont'd)
Miscellaneous information
References
Questions (and answers)

Where to find documentation? `http://bccd.net`

What software is available? `module avail`

How to load new software? `module load name`

How to unload software? `module unload name`

How to reset networking? `sudo /bin/bccd-reset-network`, use the password you set when you booted



References

- [1] Jeannette M. Wing. Computational thinking. *Commun. ACM*, 49:33–35, March 2006.

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)



Questions (and answers)

Sponsors

Preliminary work

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Live Demo! (cont'd)

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

GalaxSee - Multi-node (cont'd)

Miscellaneous information

References

Questions (and answers)