



Bootable Cluster CD Intro Parallel 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`

Sam Leeman-Munk

Shodor Foundation

`sleemanmunk at gmail.com`

Skylar Thompson

University of Washington

`skylar at cs.earlham.edu`

June 29, 2011



Sponsors

Sponsors

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

Miscellaneous information

References

Questions (and answers)

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



What is the BCCD?

Sponsors

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

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

What is the BC CD?

Short history

What is the problem?

How is BC CD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

Miscellaneous information

References

Questions (and answers)

- First versions developed by Paul Gray and Students at University of Northern Iowa
- Current version (BC CDv3) collaboratively developed by a number of educational institutions
- BC CDv3 much easier to maintain and able to respond to new needs
- BC CDv3 has the same user experience as BC CDv2, so curricula written for v2 will still work with v3



What is the problem?

Sponsors

What is the BC CD?

Short history

What is the problem?

How is BC CD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

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

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

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

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

Miscellaneous information

References

Questions (and answers)

- Double click on VMware player
- Go to File-> New
- Click on “Install OS Later”
- Select “Linux”, then make sure “Ubuntu” is selected
- Set the name to “BCCD”
- Take the defaults for the storage
- Click Customize
- Set Procs-> 4 (if your system says QUAD) or 2
- Set the Network Adapter-> Bridged
- **Don't boot yet!**



Boot

Sponsors

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

Miscellaneous information

References

Questions (and answers)

- Just keep pressing Enter.

GalaxSee - Single node

Sponsors

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

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

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

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 VMware, because of virtualization overhead.

GalaxSee - Multi-node

Sponsors

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

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 8 \  
/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

What is the BC CD?

Short history

What is the problem?

How is BC CD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

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

Sponsors

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

Miscellaneous information

References

Questions (and answers)

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



Questions (and answers)

Sponsors

What is the BCCD?

Short history

What is the problem?

How is BCCD a solution?

Live Demo!

Boot

GalaxSee - Single node

GalaxSee - Single node (cont'd)

GalaxSee - Multi-node

Miscellaneous information

References

Questions (and answers)