

Kittyhawk

Boston University SESA Group
Dan Schatzberg


Blue Gene/P

System



72 Racks
73728 nodes
294912 cores
1 PF
144 TB DDR2


Rack



μ


32 node-cards
1024 nodes
4096 cores
13.9 TF
2 TB DDR2

Node
Card



32 nodes
128 cores
435 GF
64GB DDR2

Compute
Card

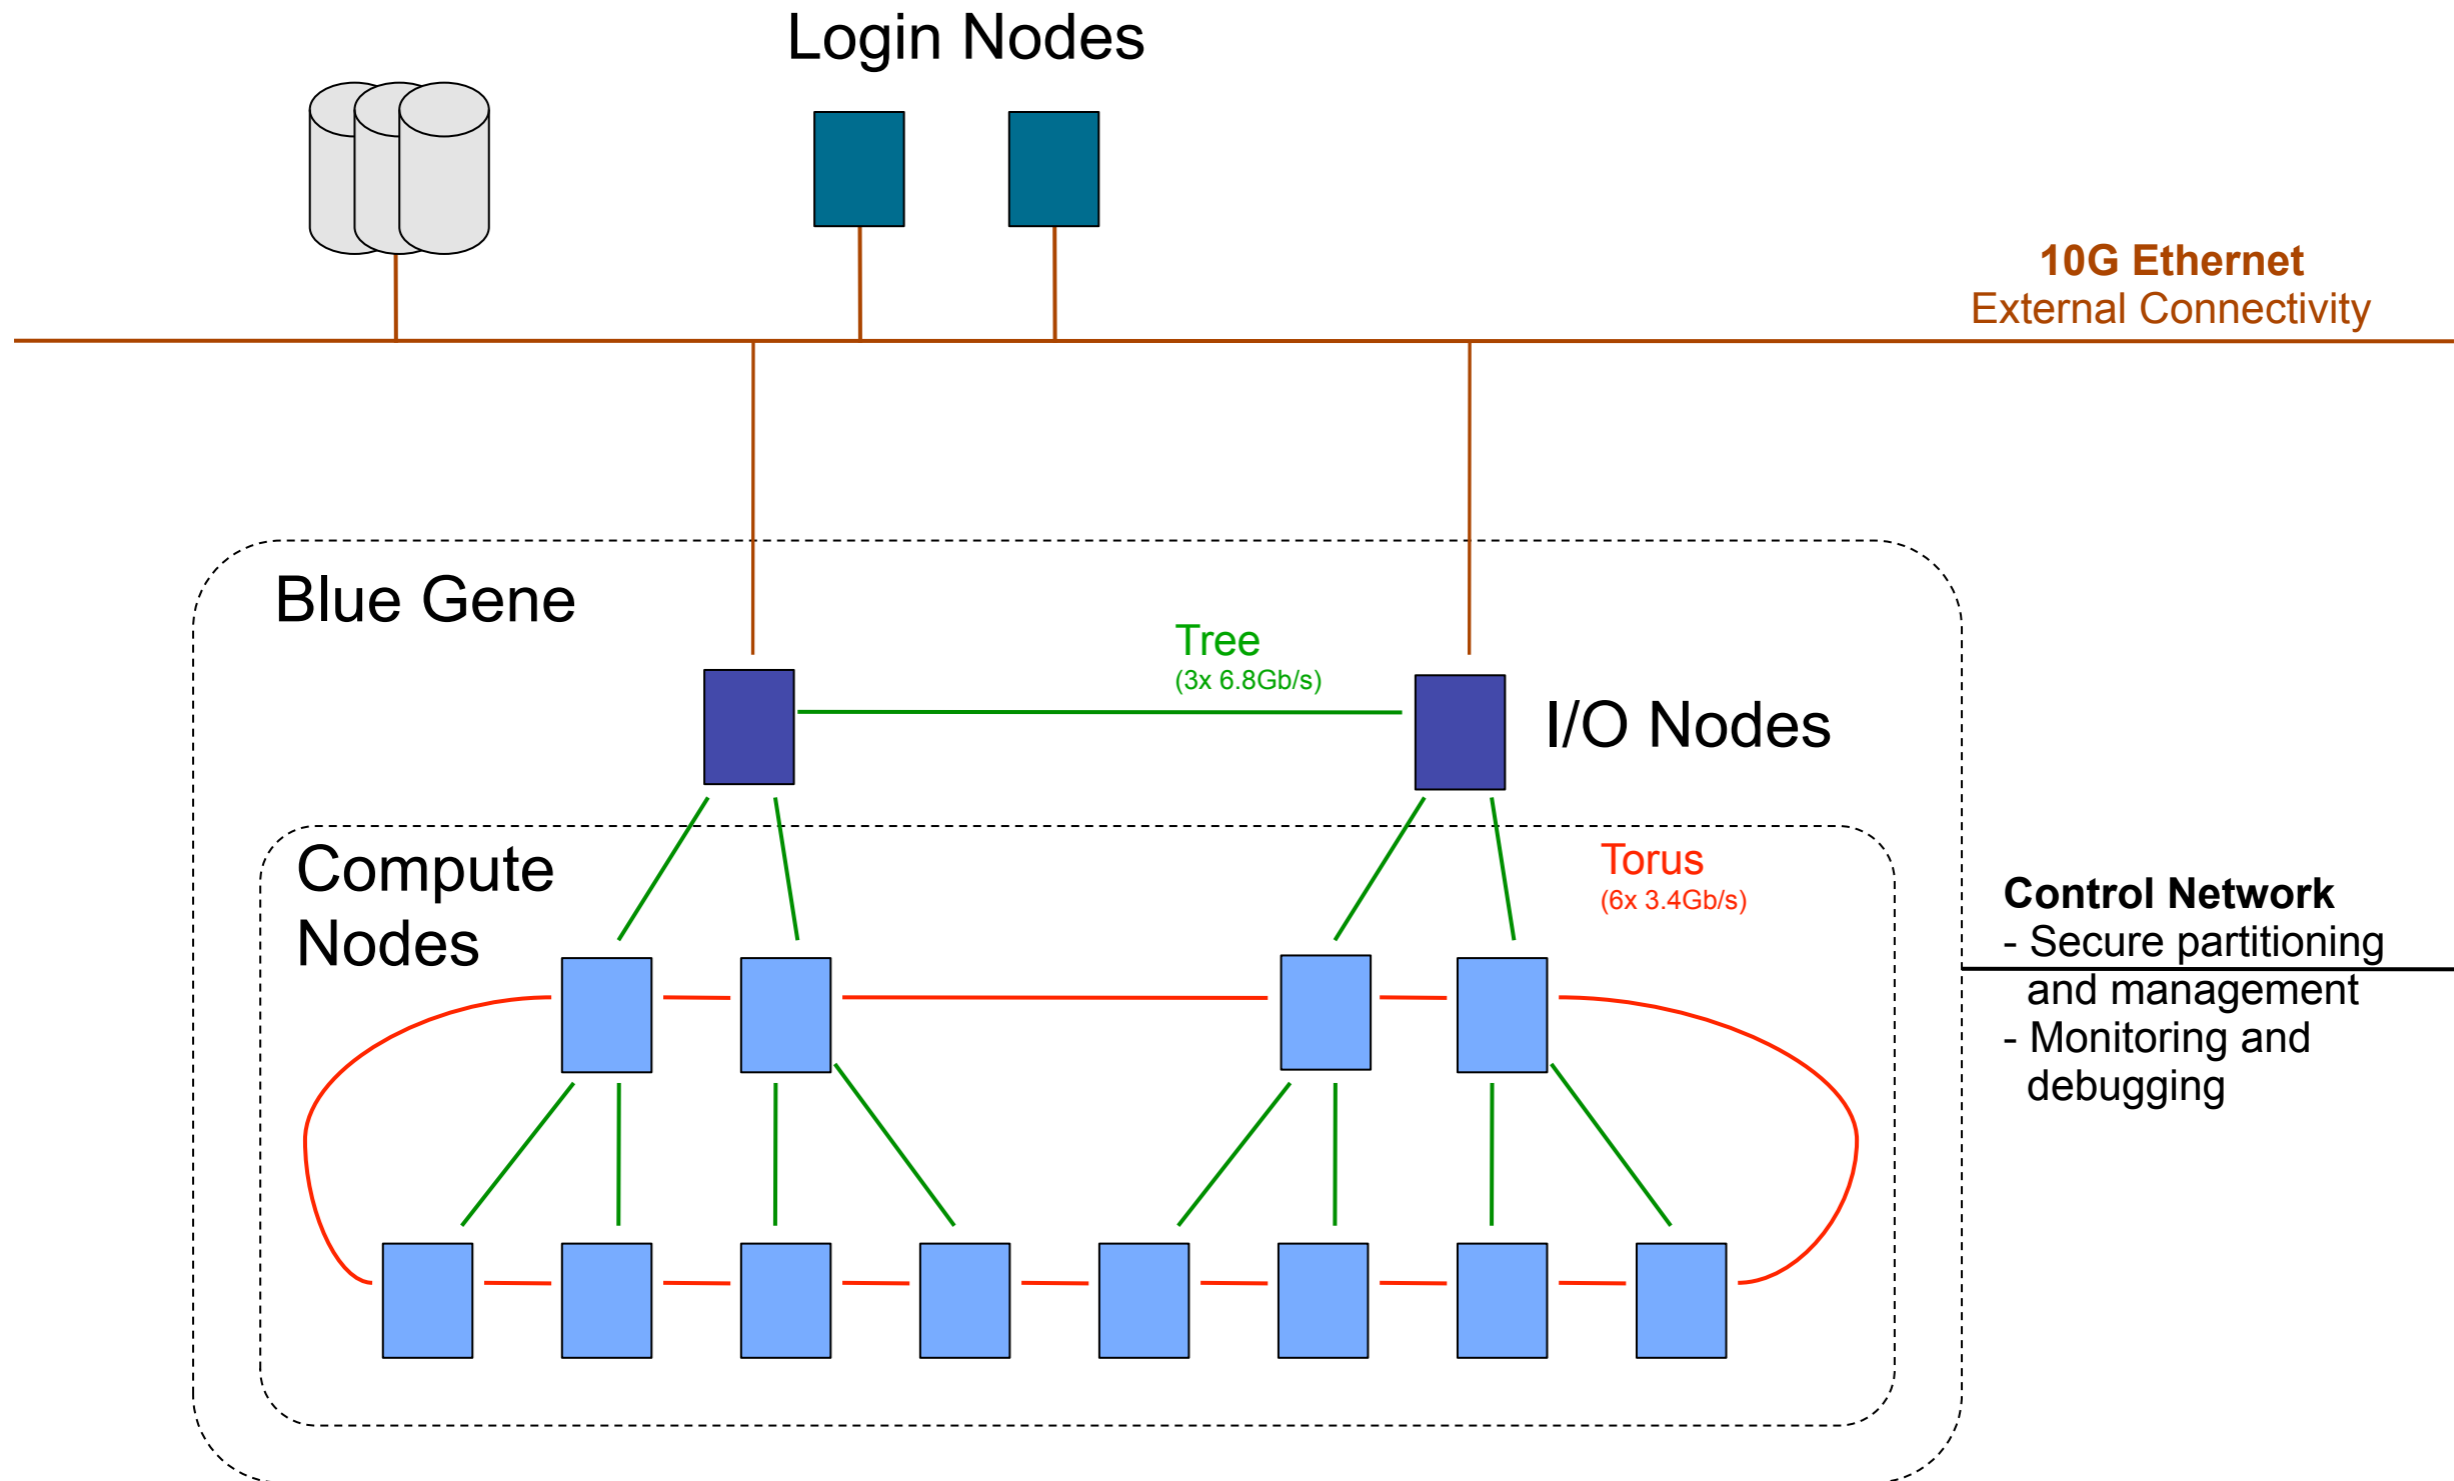


1 node
4 cores
13.6 GF
2GB DDR2

Chip



4 cores
13.6 GF



10G Ethernet
External Connectivity

Blue Gene

Tree
(3x 6.8Gb/s)

I/O Nodes

Compute
Nodes

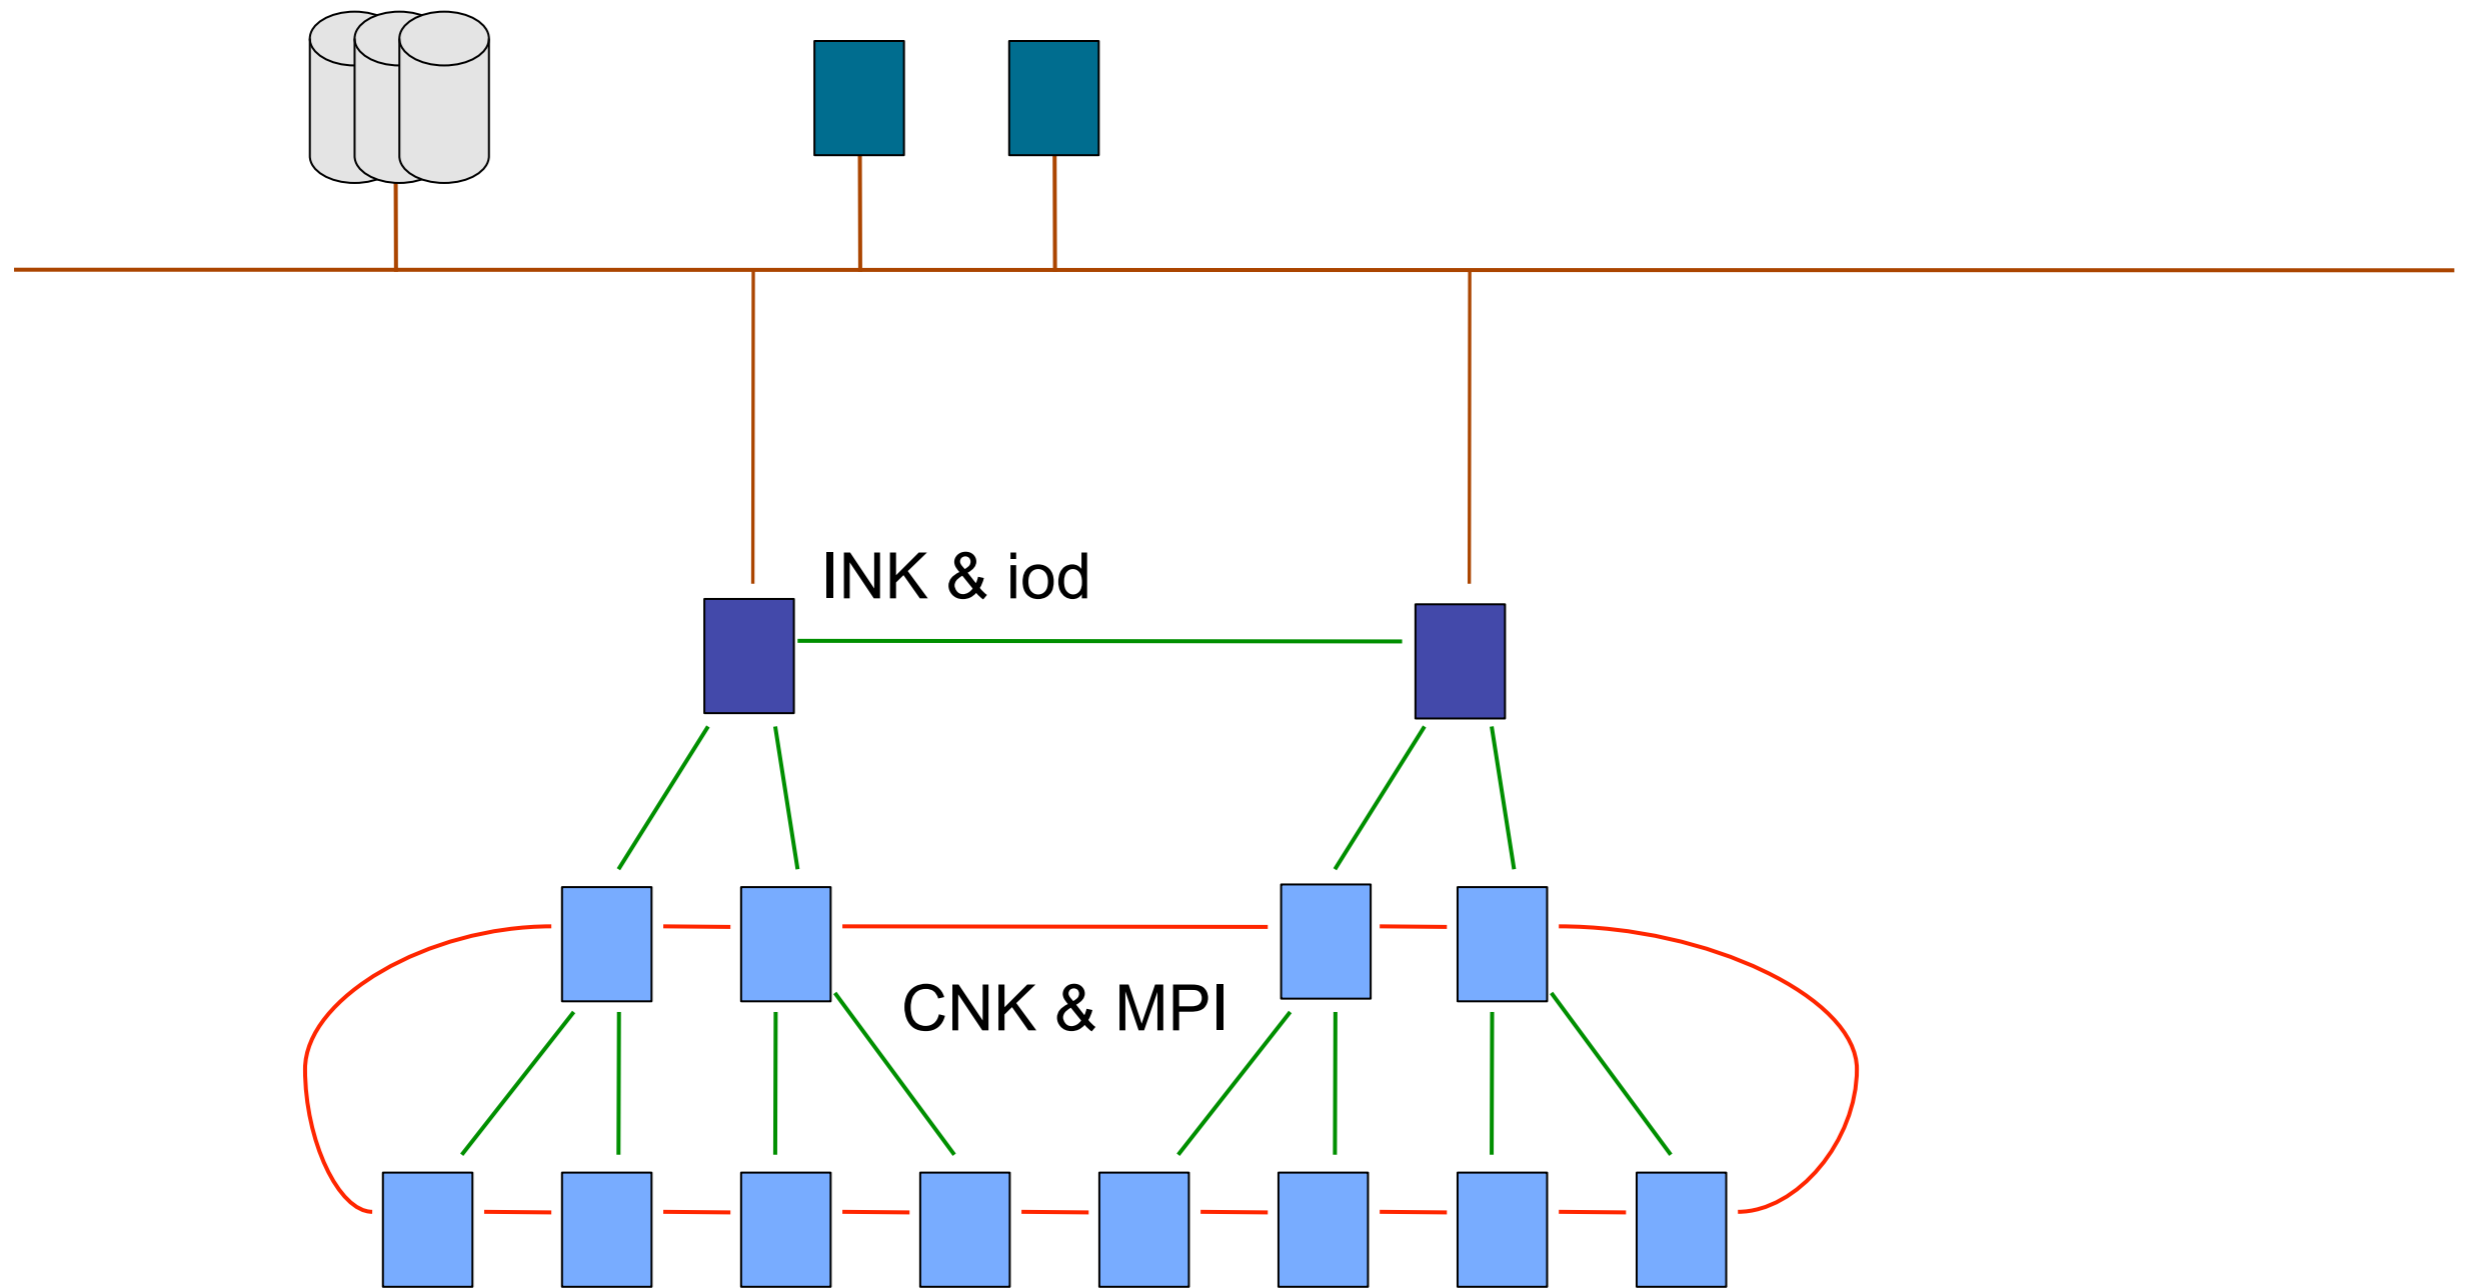
Torus
(6x 3.4Gb/s)

Control Network
- Secure partitioning
and management
- Monitoring and
debugging

Kittyhawk

- Dynamically allocate nodes to many users
- Allow arbitrary network configurations
- Allow arbitrary system software

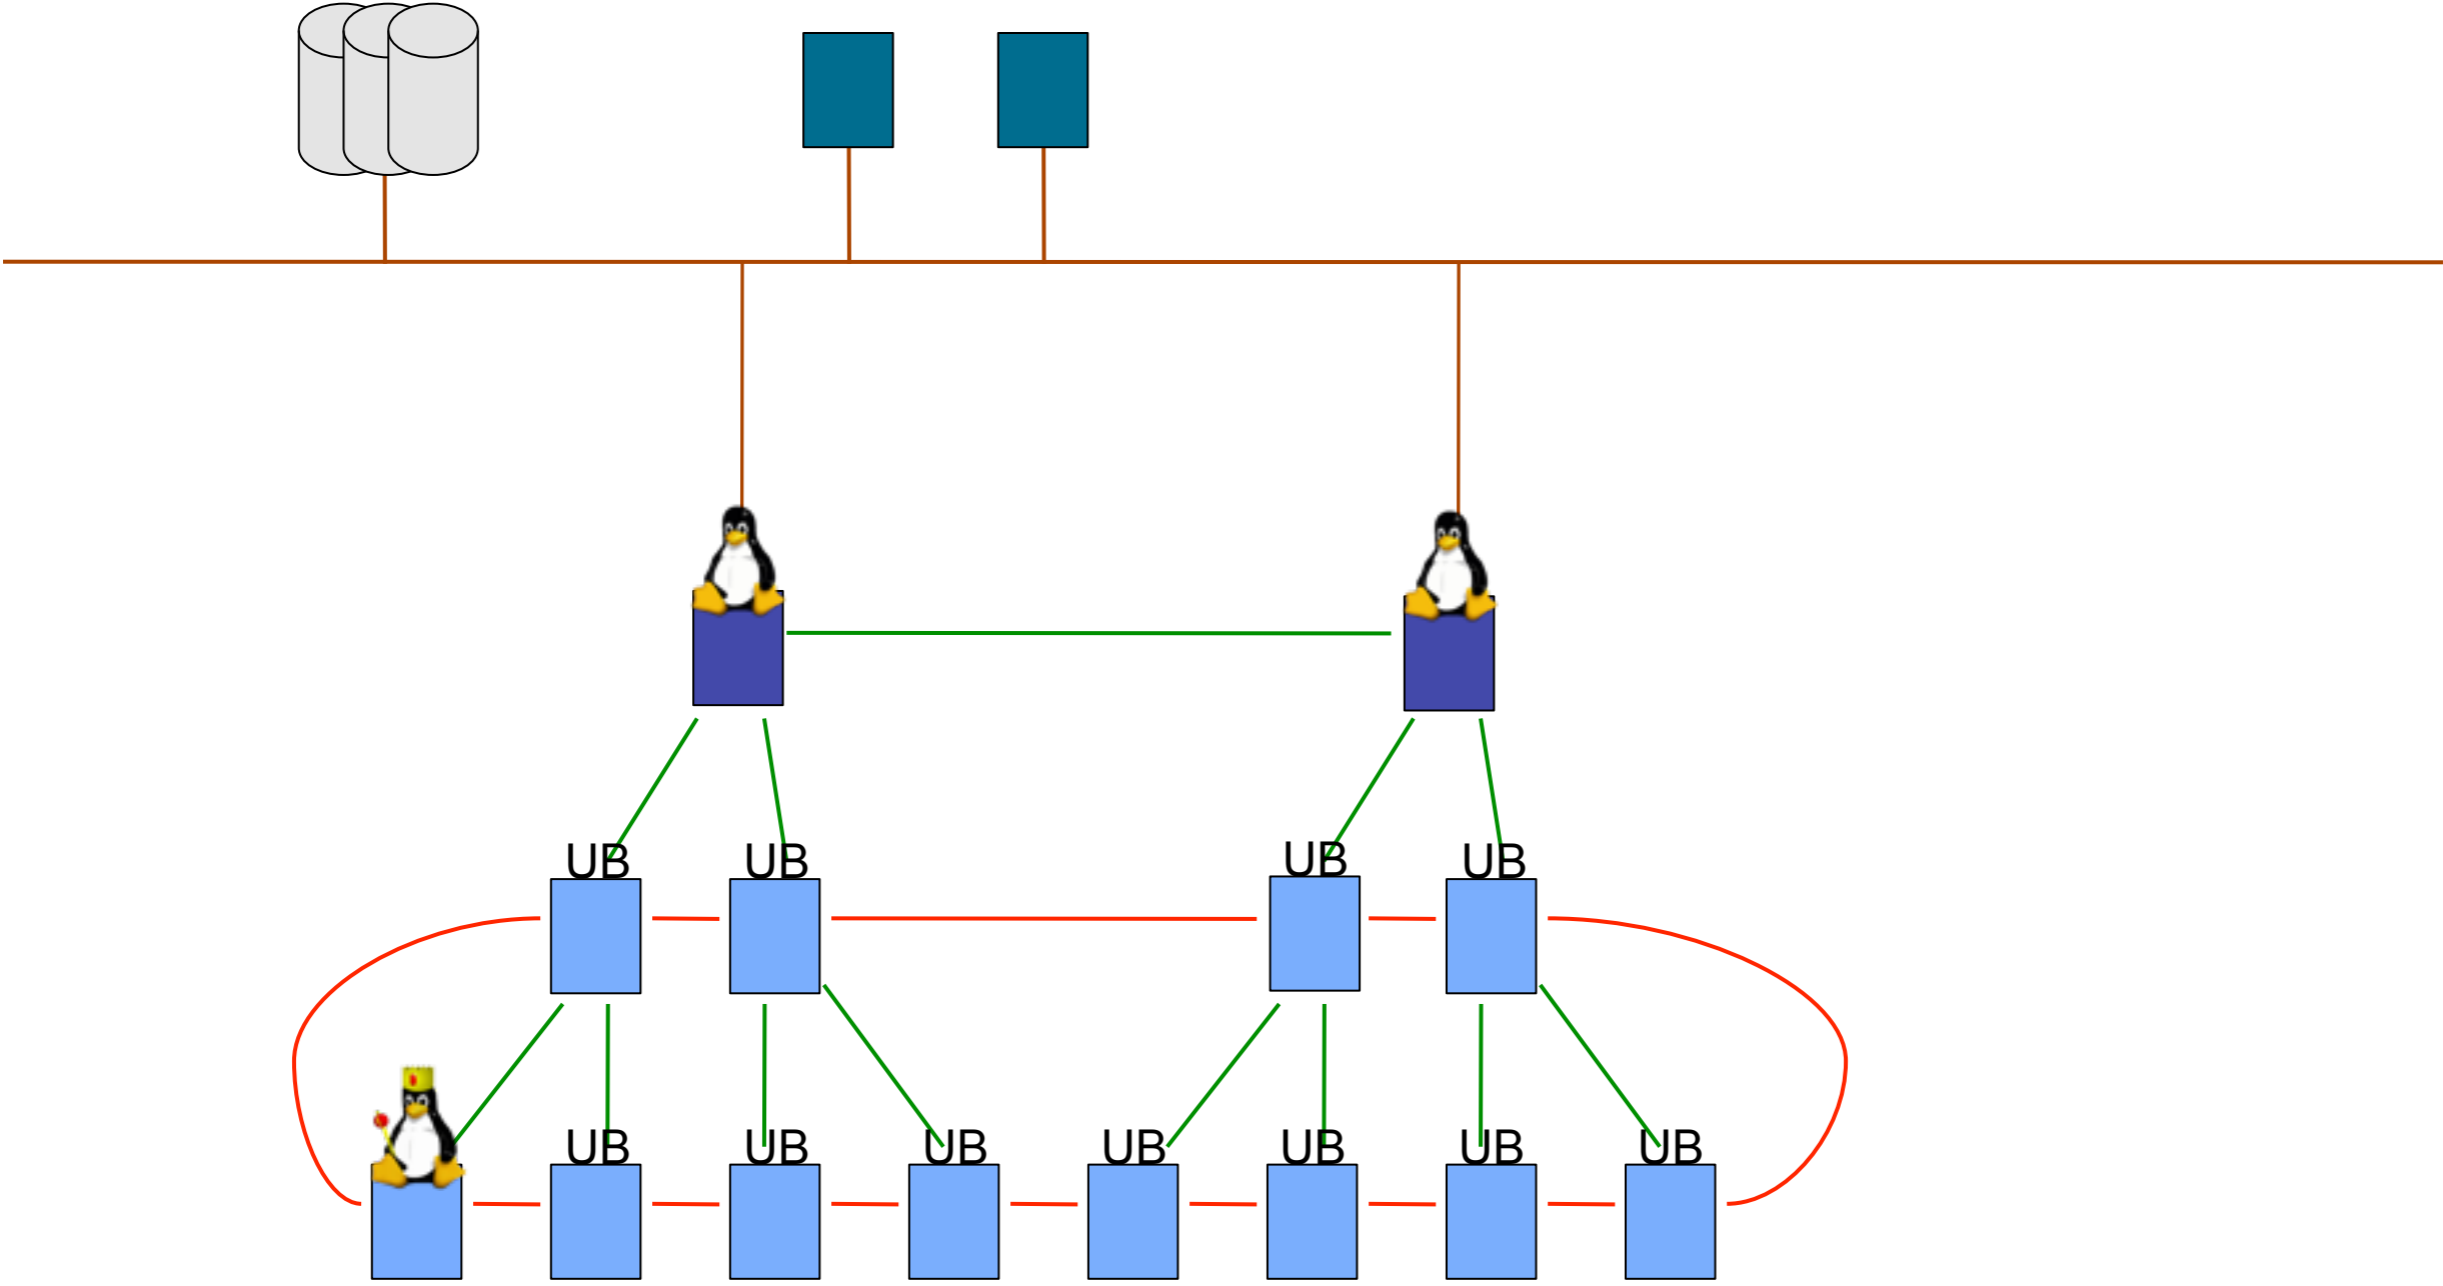
```
qsub -n <nodes> -t <time> -q <queue> -A <project> --kernel  
<profile> MPIapp.elf
```



/bgsys/argonne-utils/profiles/<profile>/*

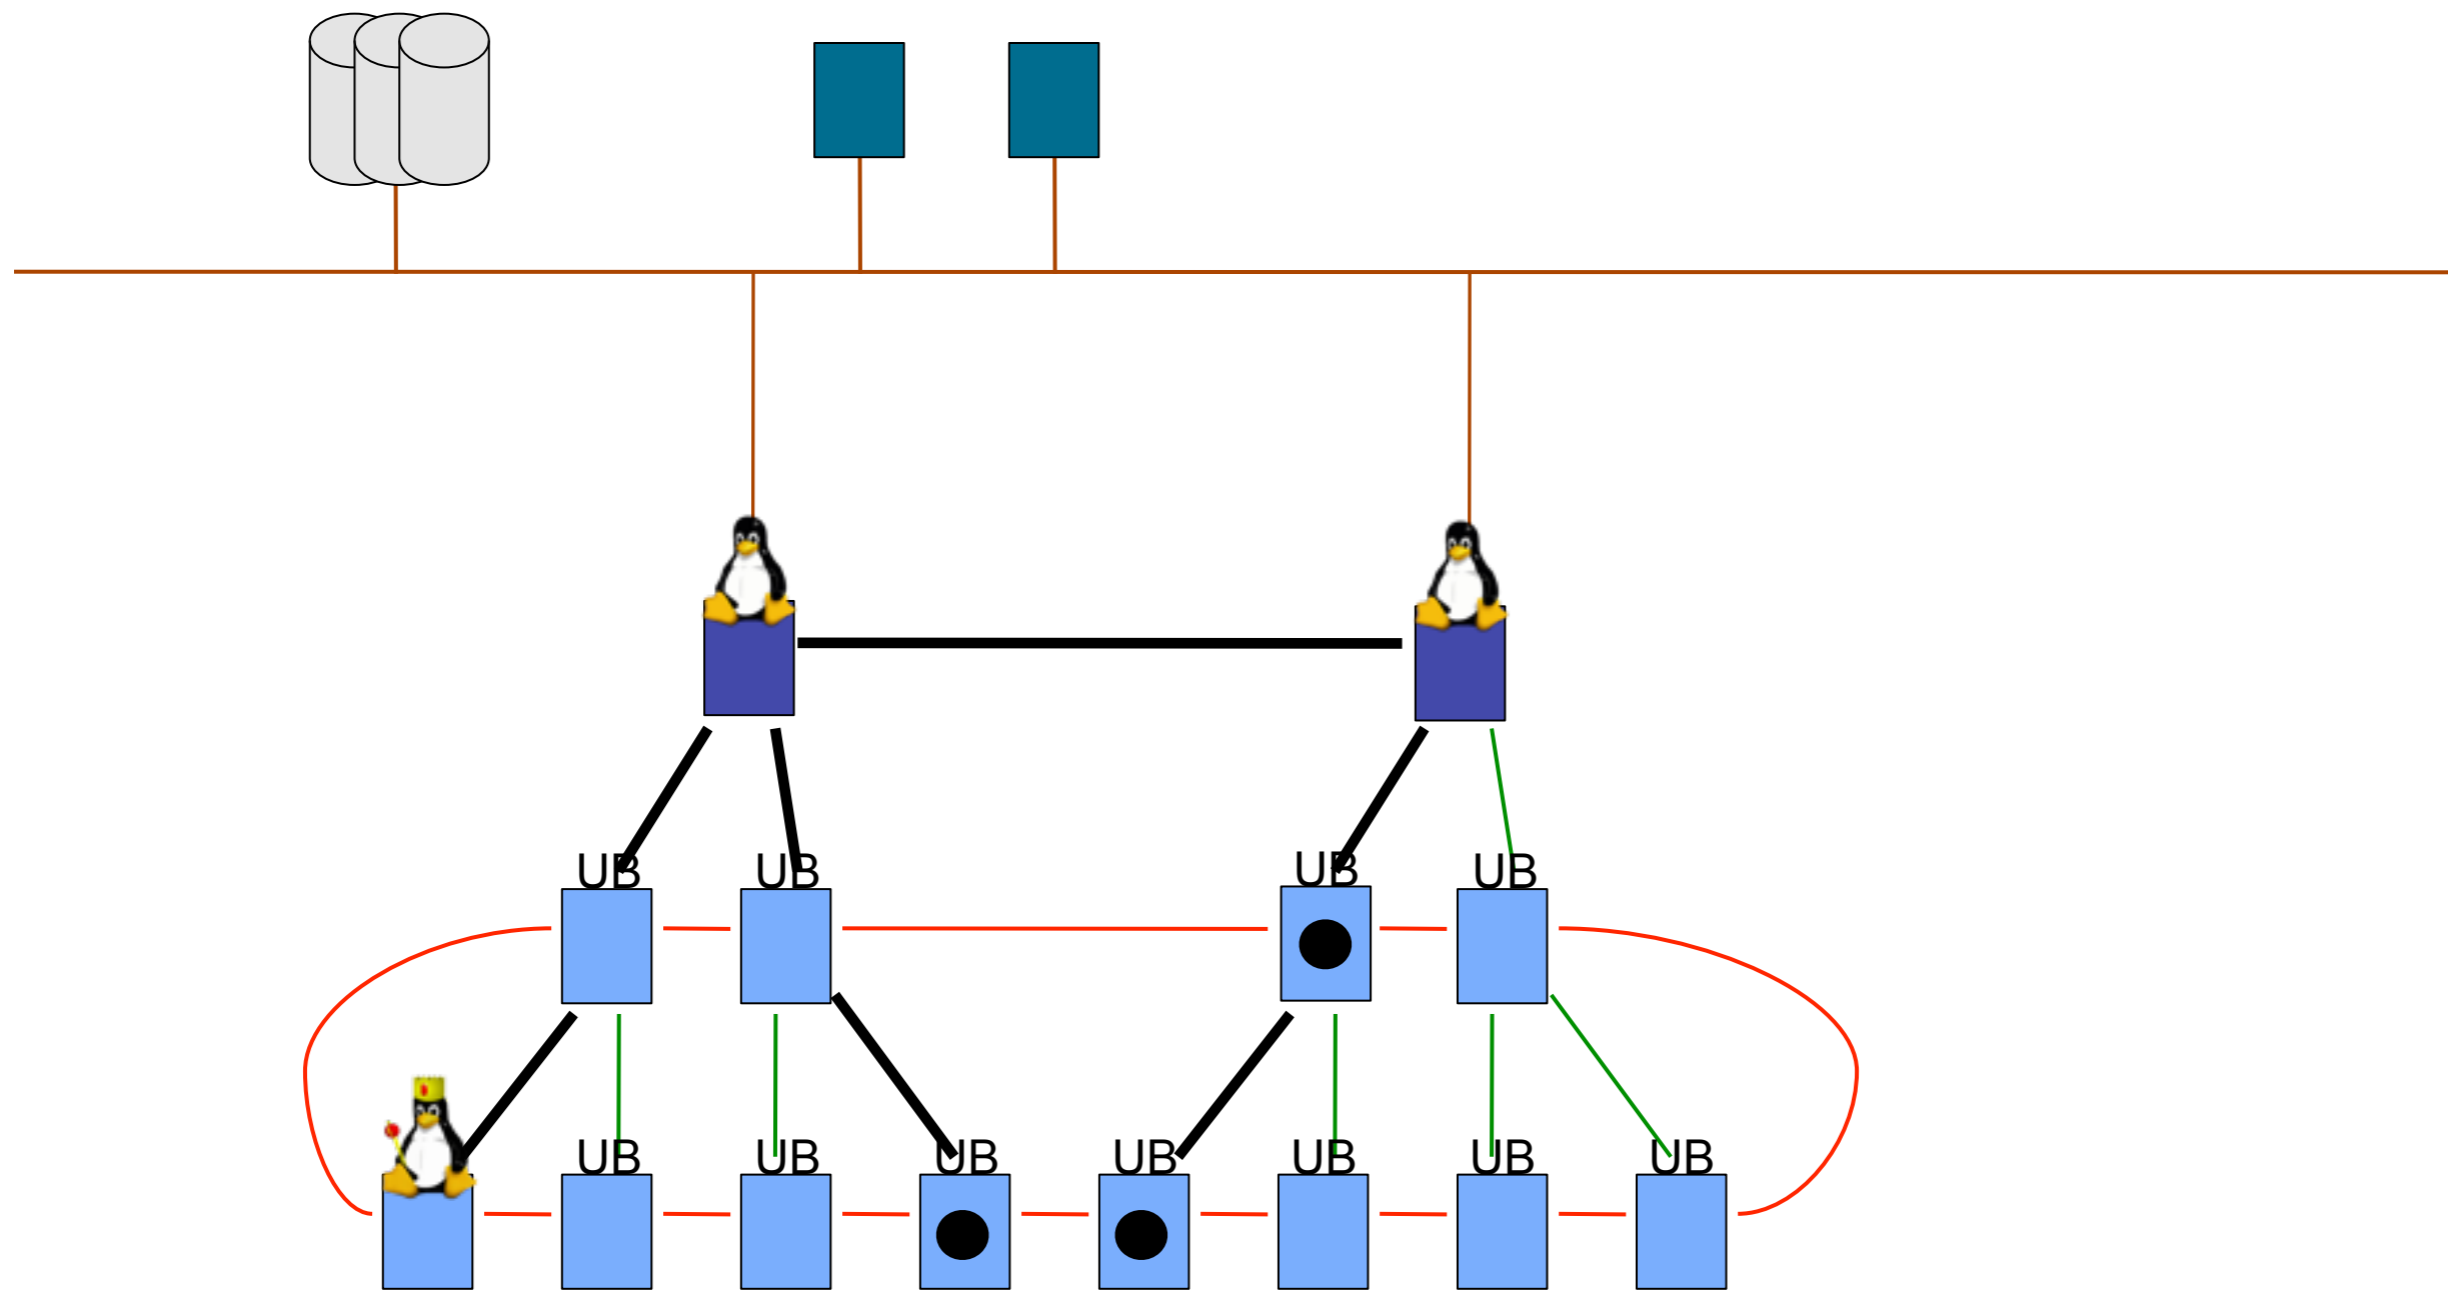
khqsub <time> <nodes>

```
qsub -n <nodes> -t <time> -e stderr -o stdout --debuglog debuglog -q default -A  
BGPPlan9Meas_DD --kernel kh $pwd/null.elf
```

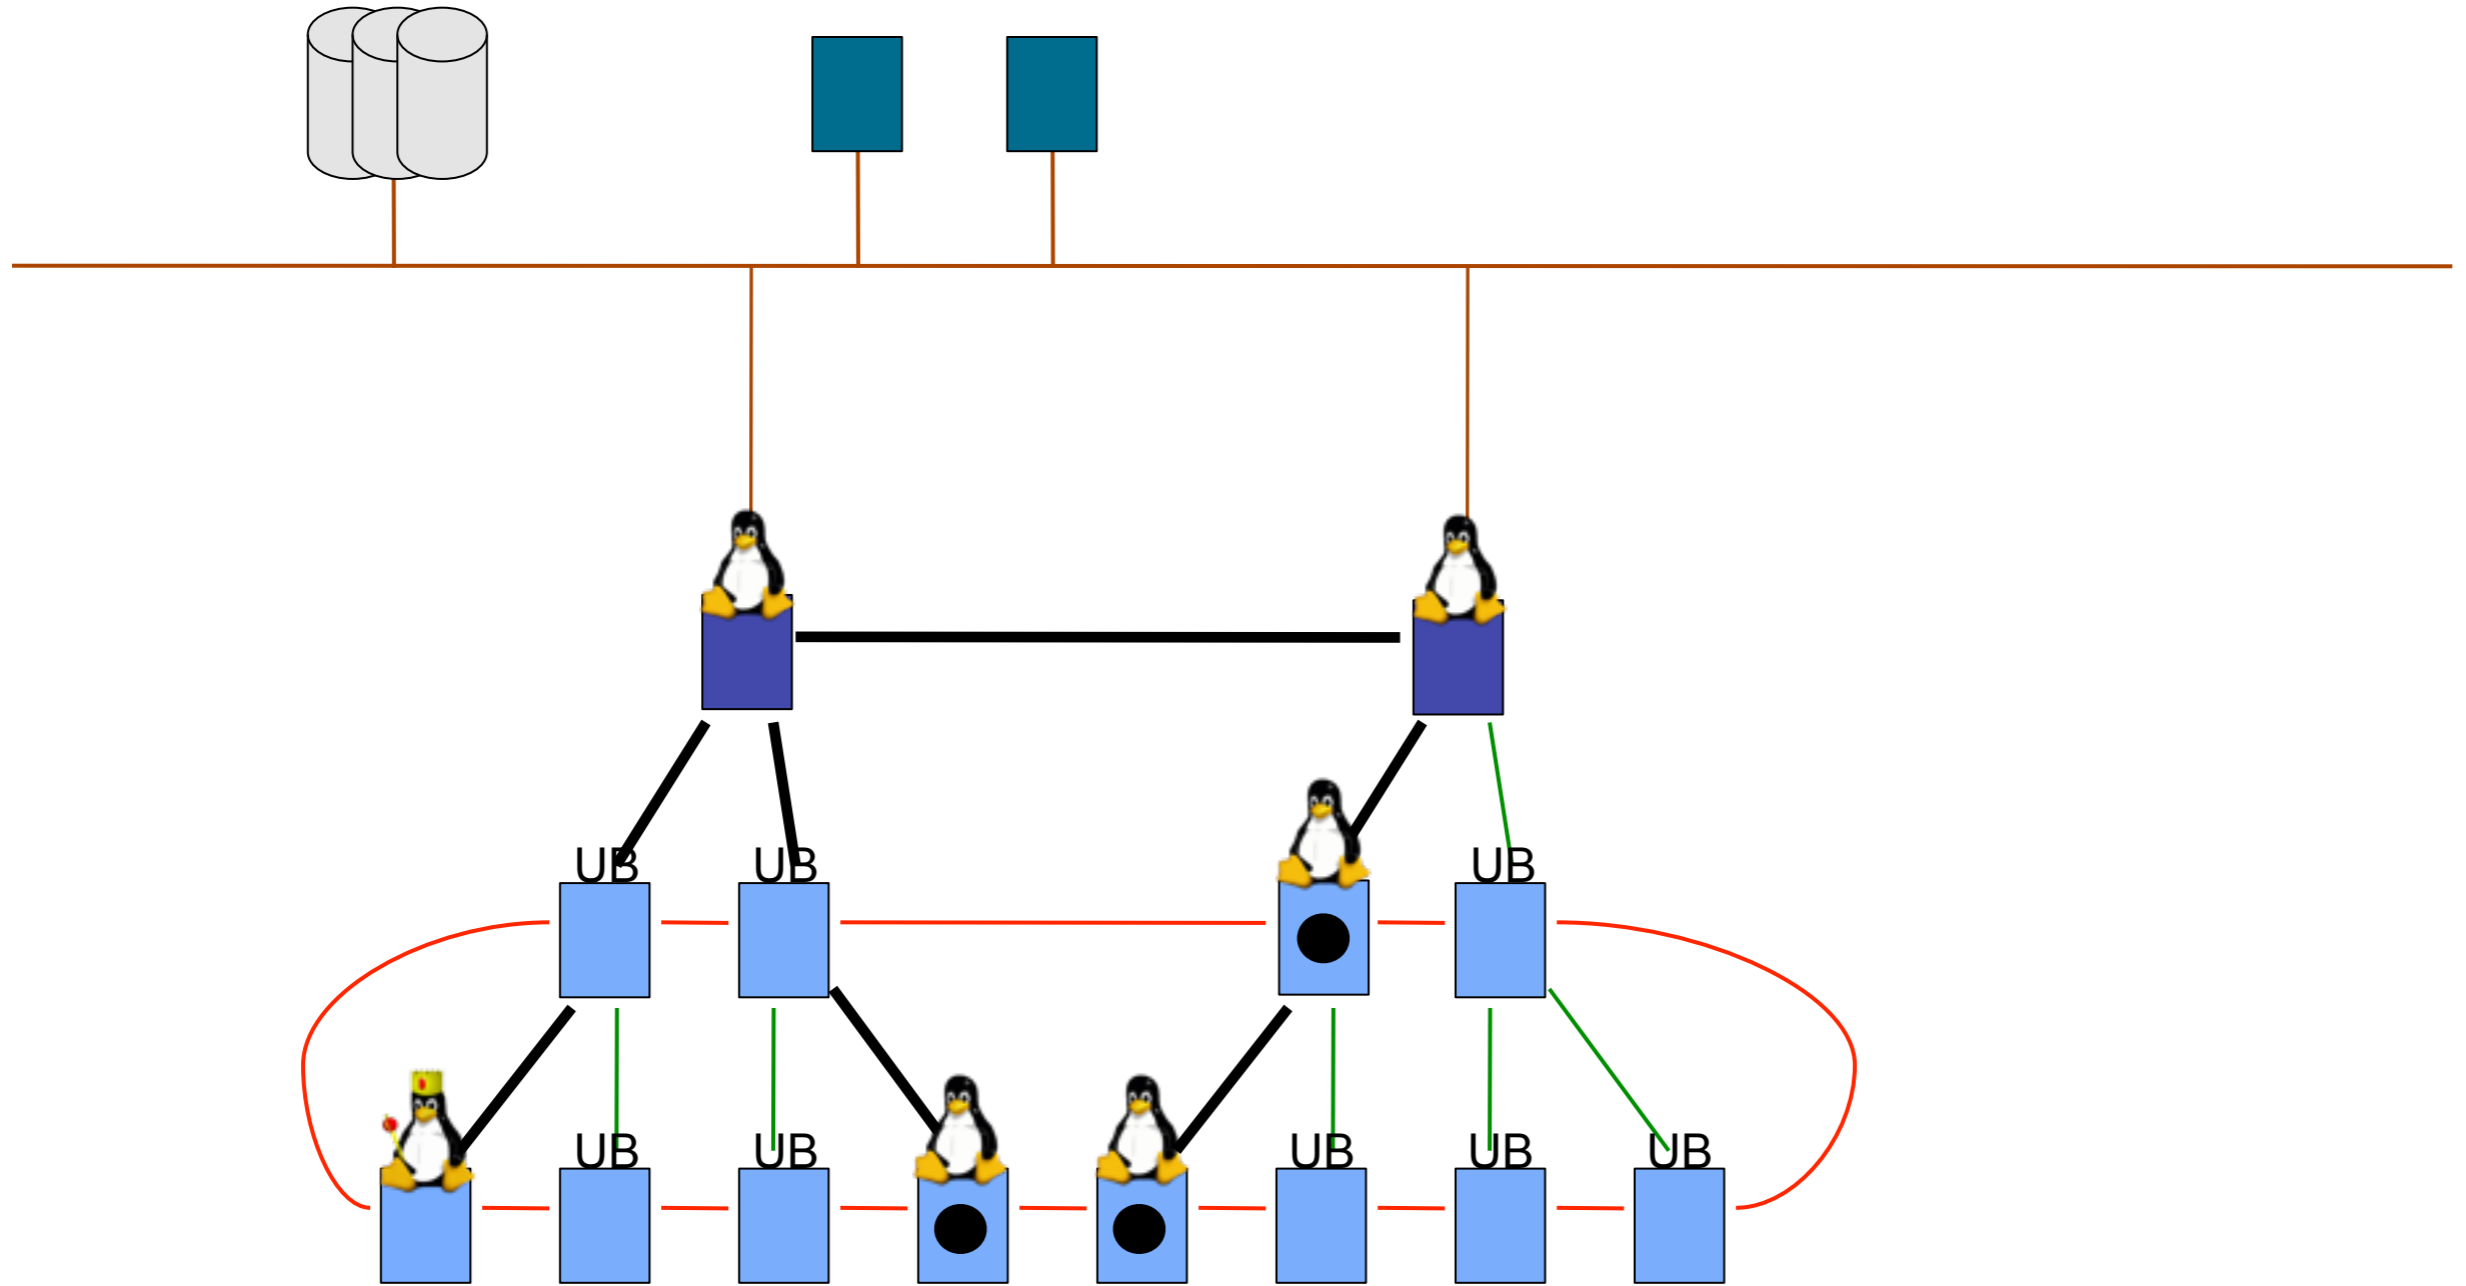


kh control server

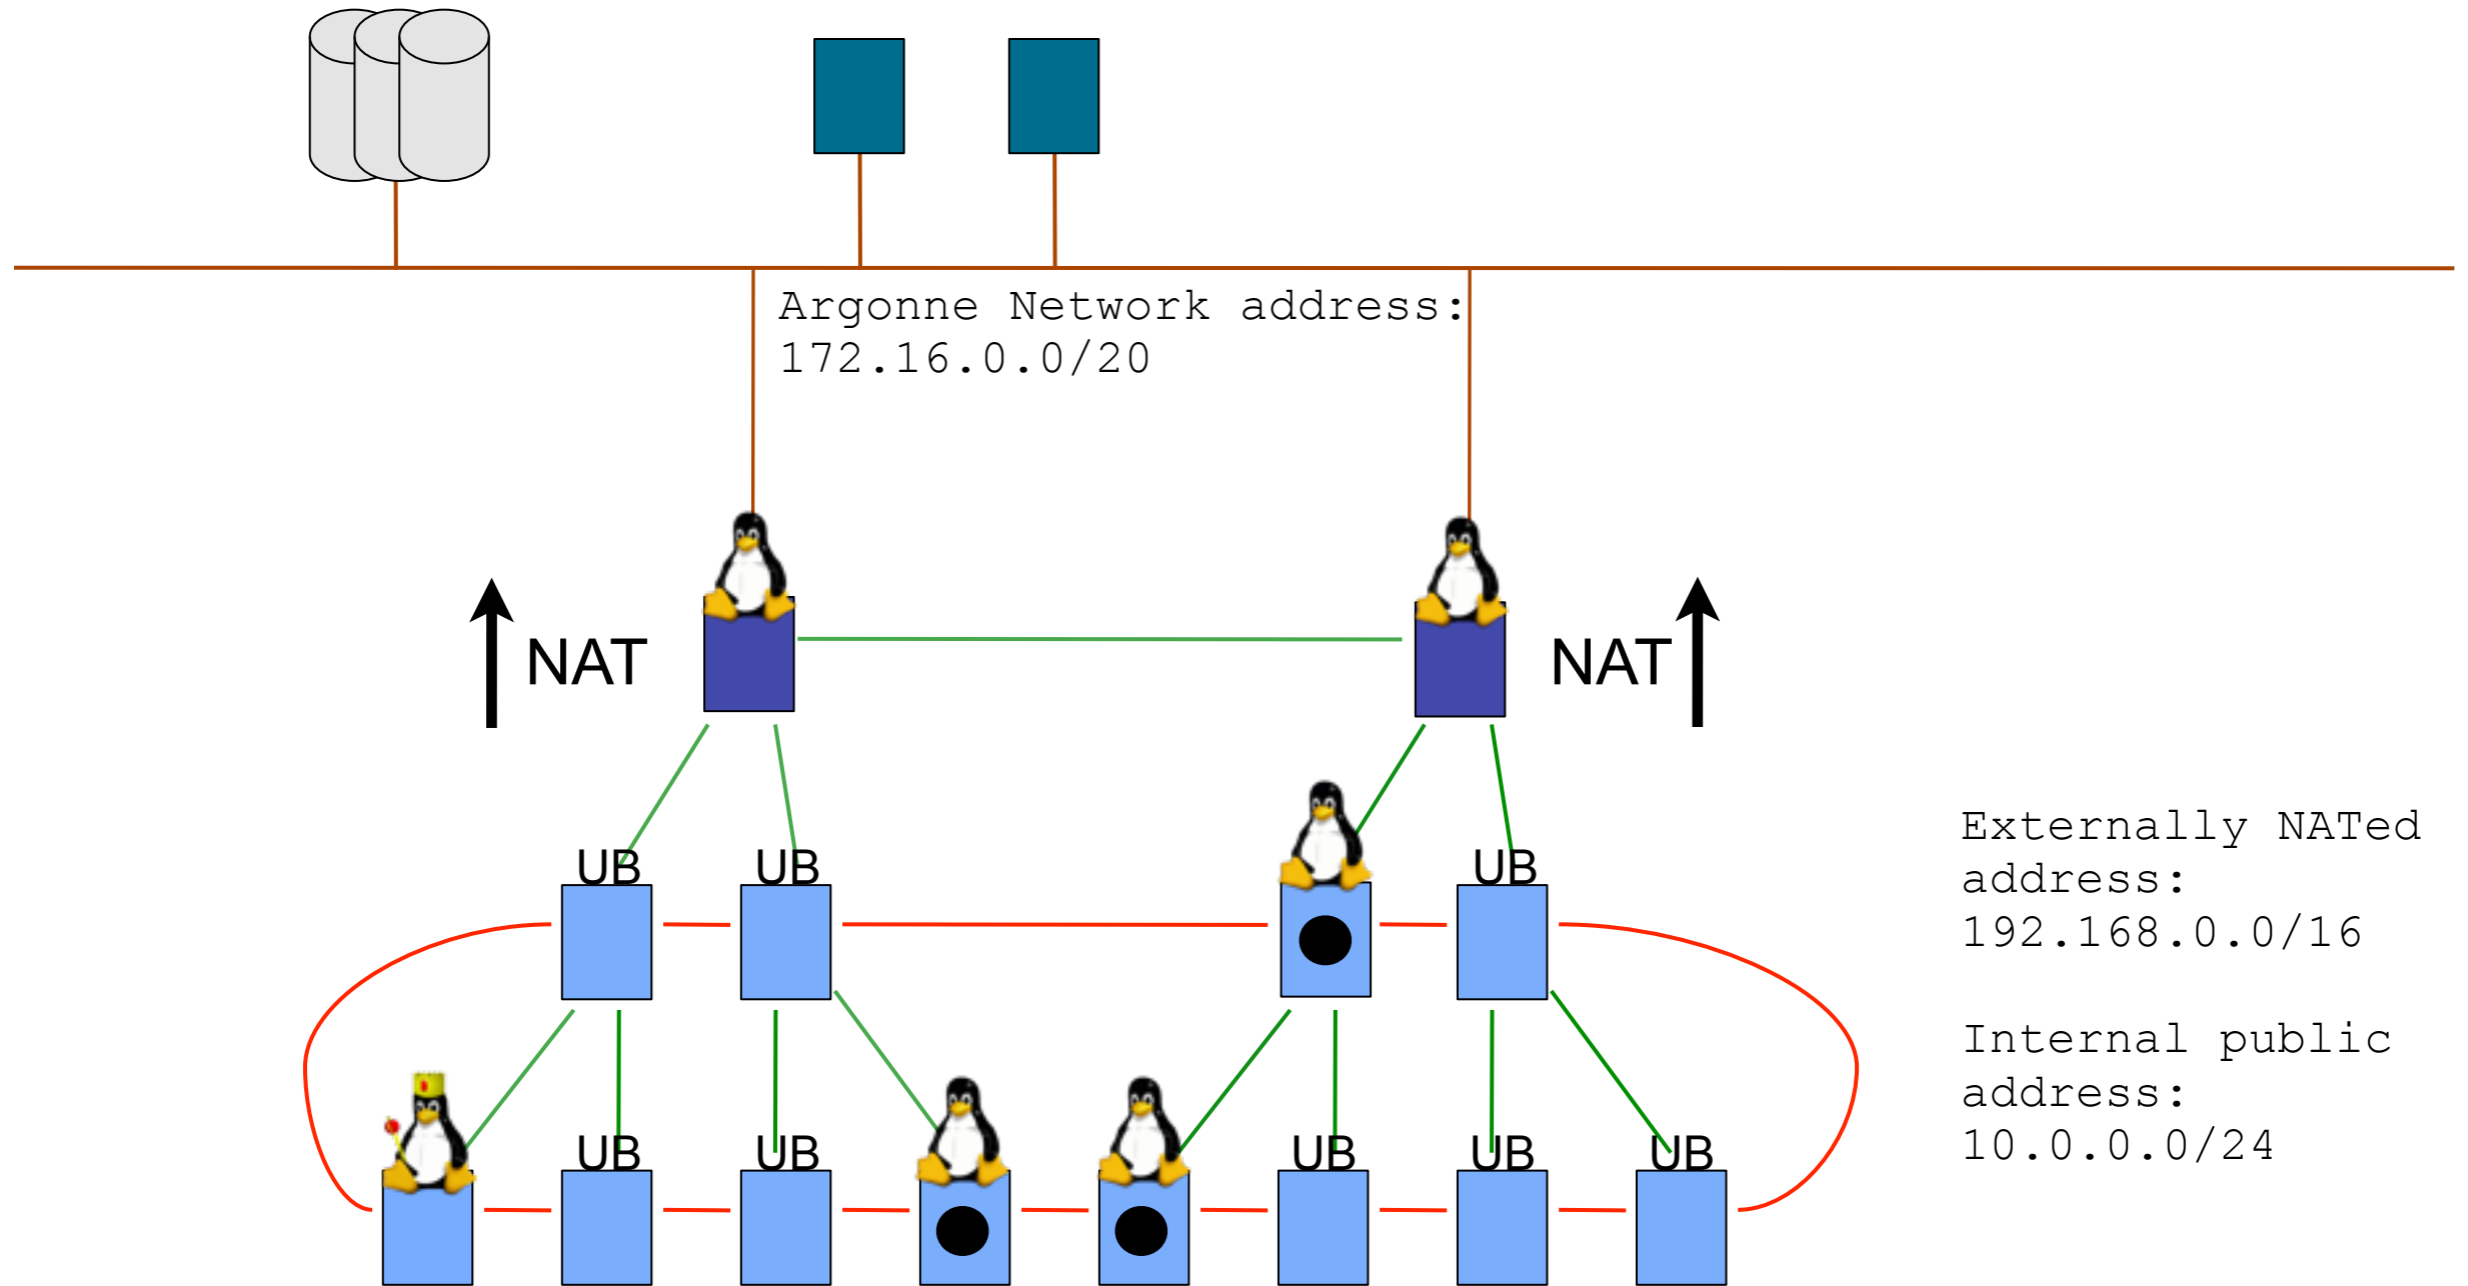
khget <user> <nodes>



```
khdo loadramdisk <ramdisk>  
khdo loadkernel <kernel>  
khdo writecon "run kboot"
```

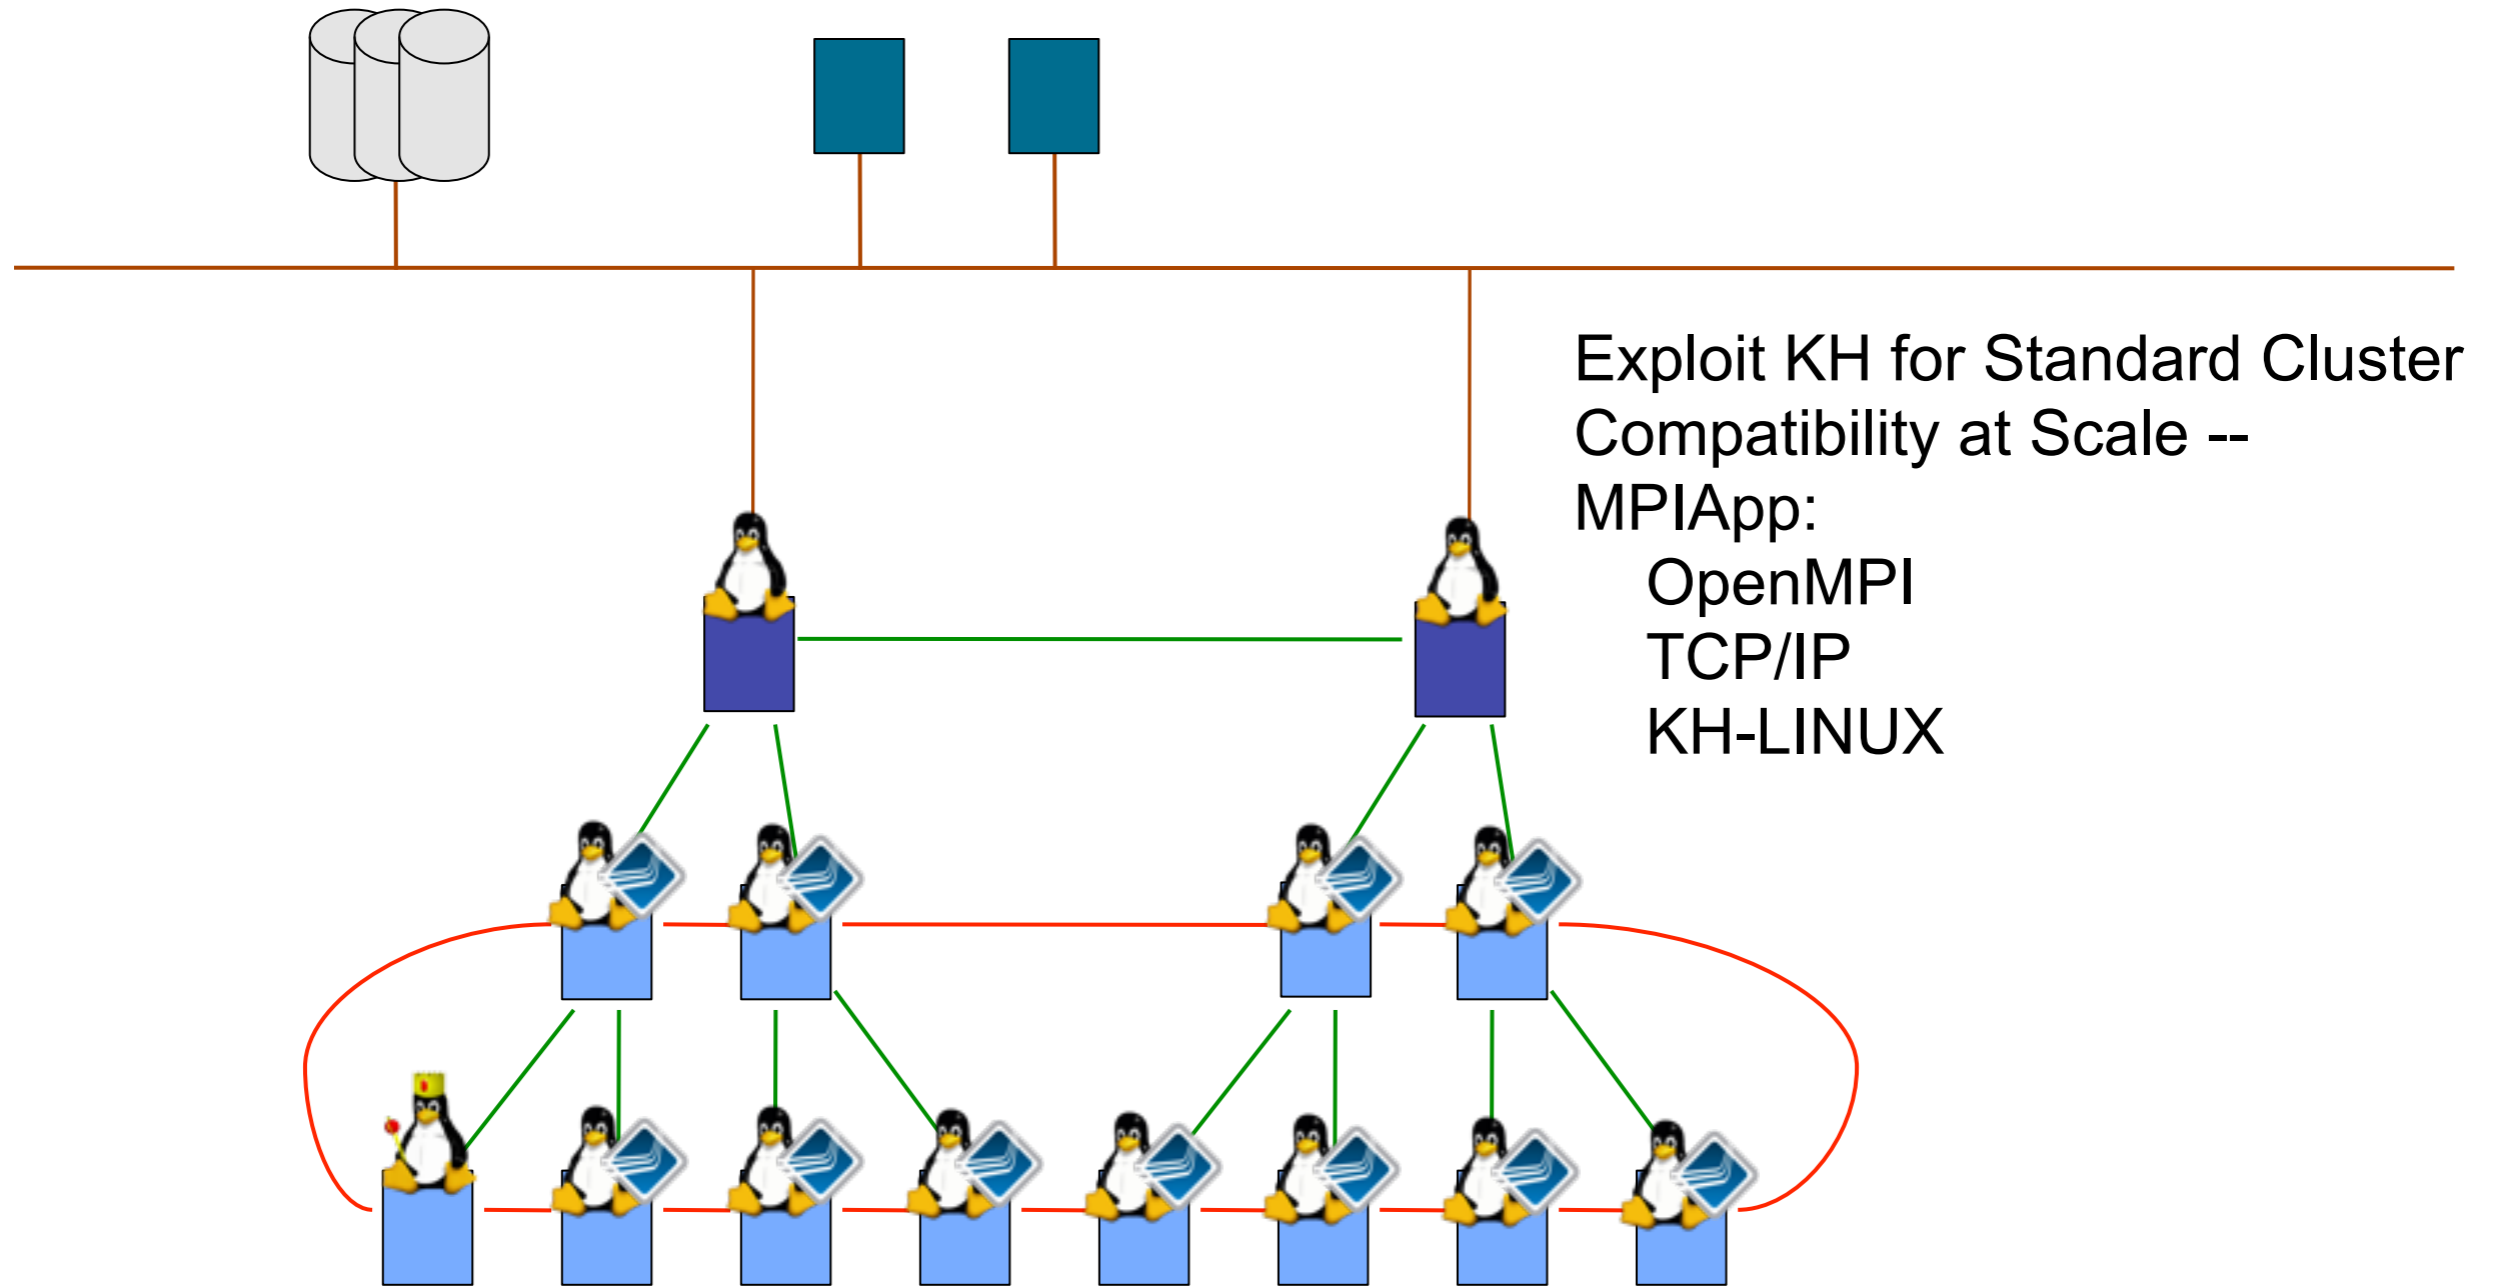


Challenger/Intrepid IP configuration

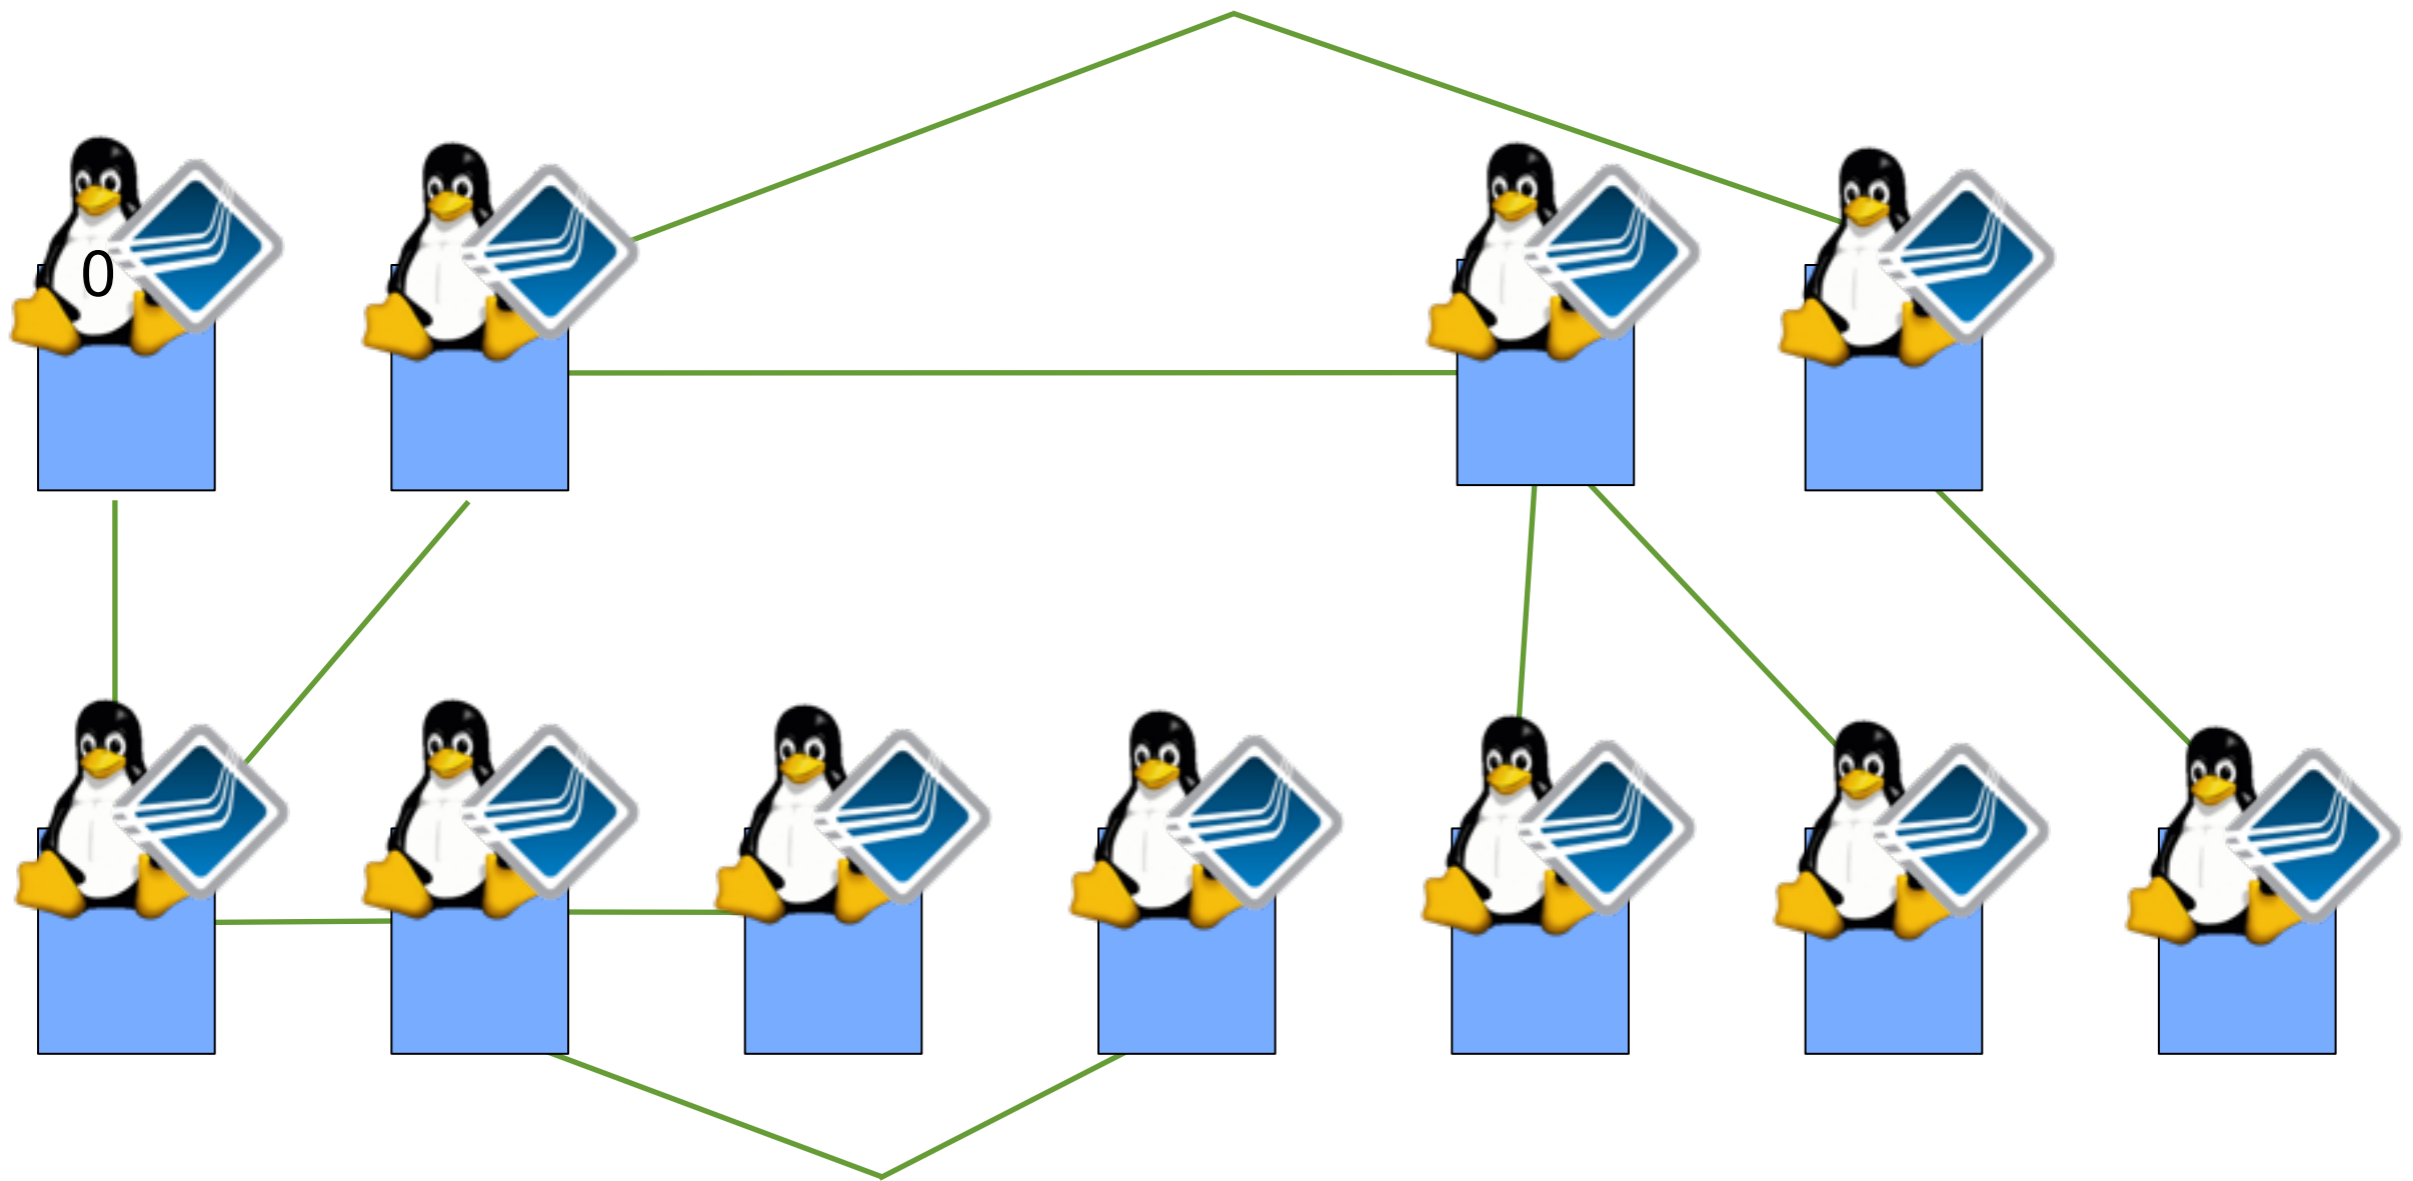


Contribution I: Kittyhawk as a fault testbed

Kittyhawk FOX Fault TestBed



OpenMPI OverLay Network (ontop of TCP/IP)

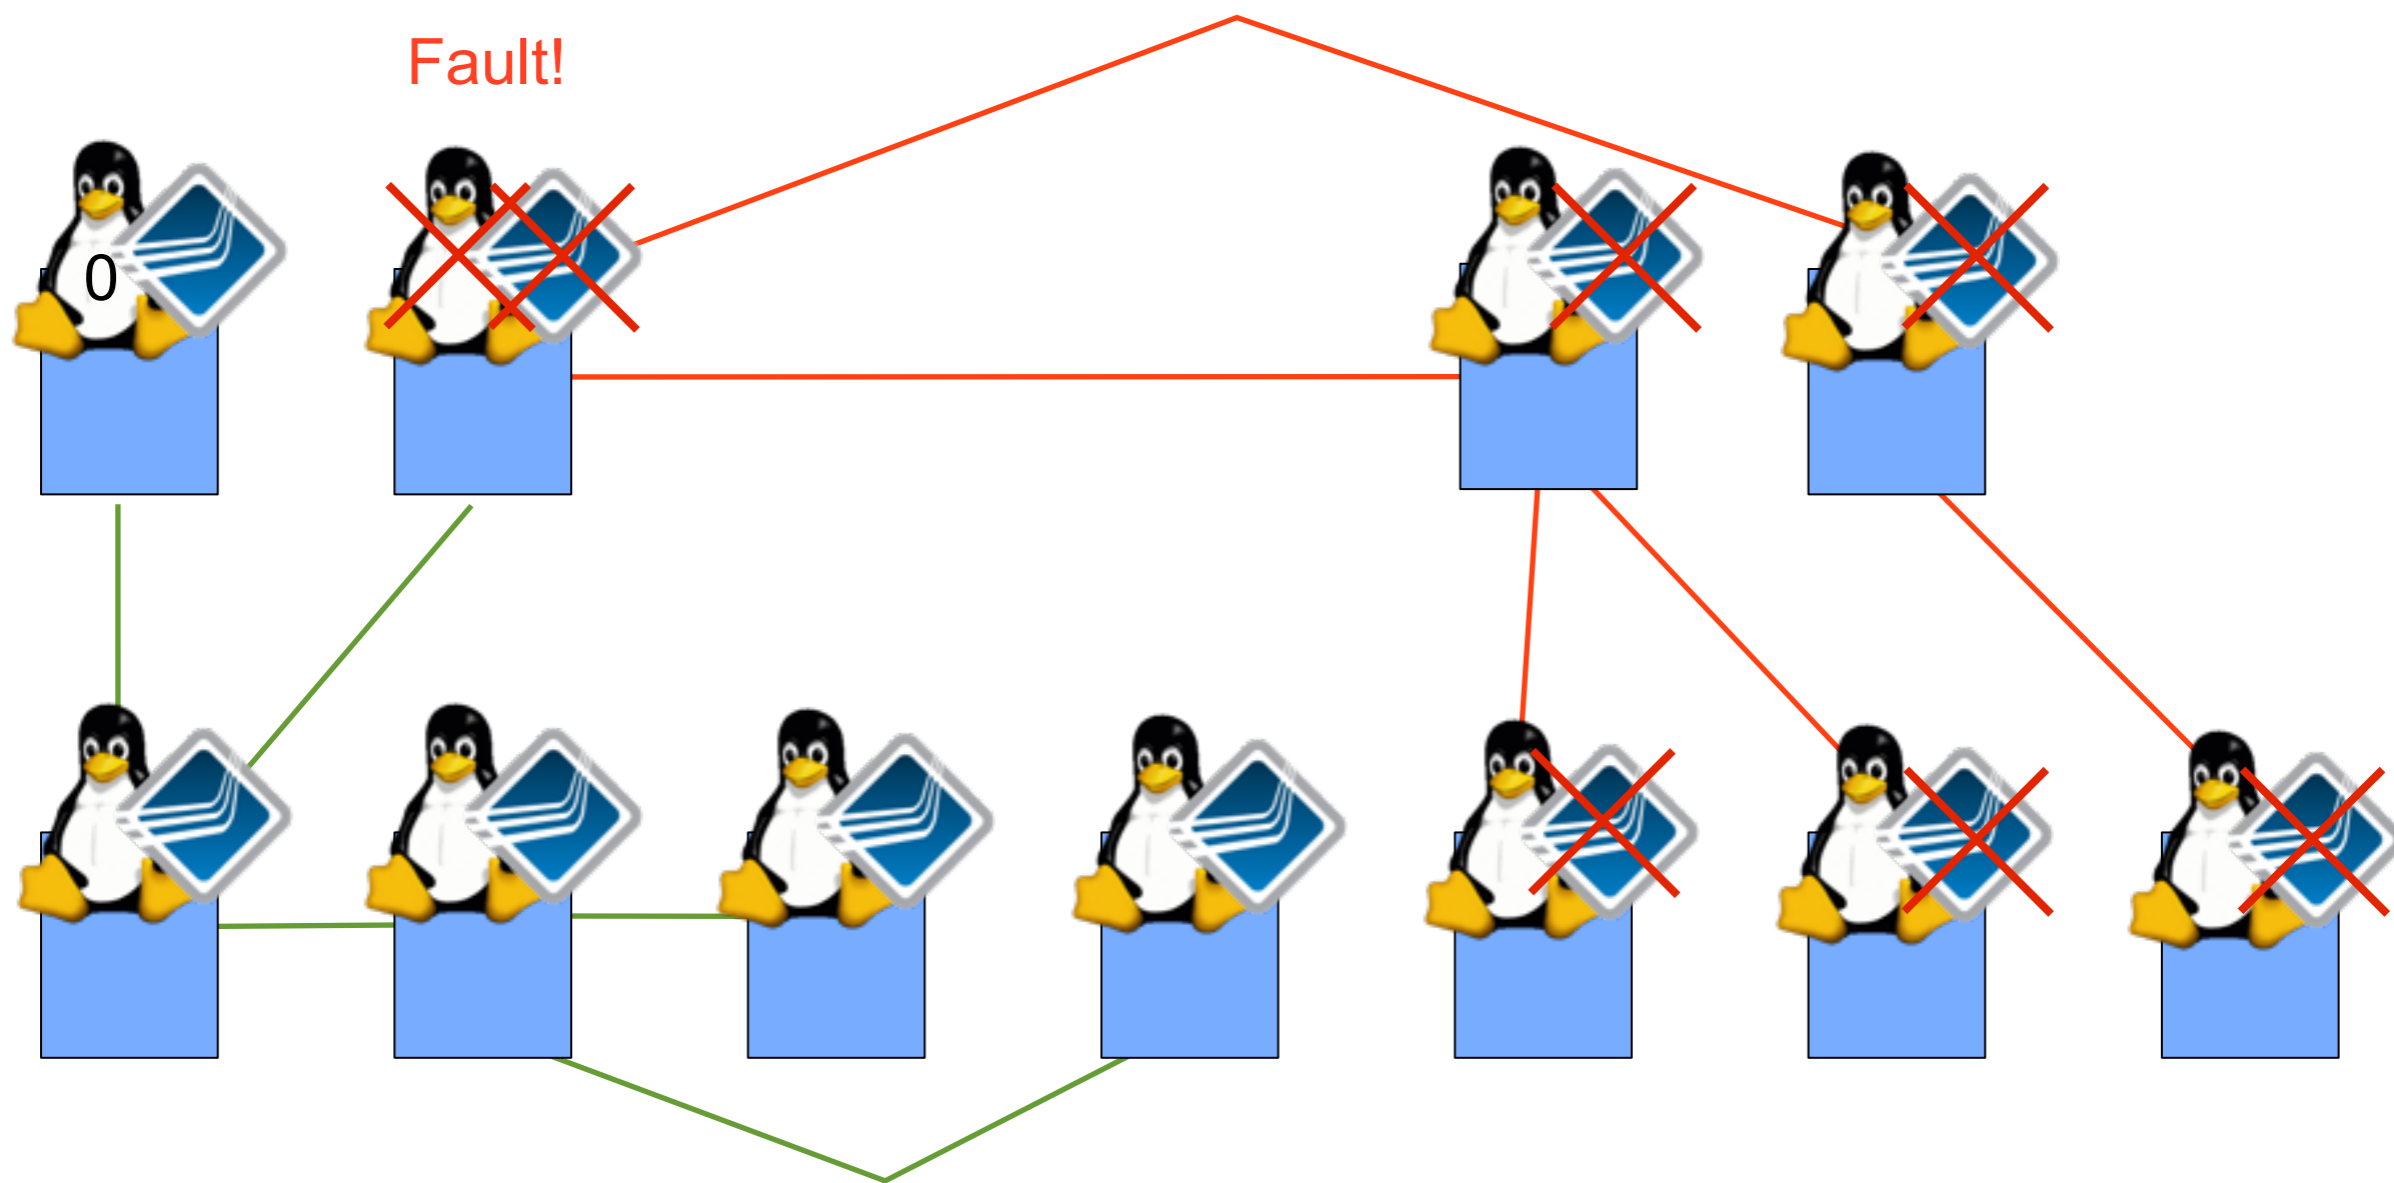


Rank -> IP Address

Internal public address:
10.0.0.0/24

Inject "Faults" by "Killing" OS Kernel

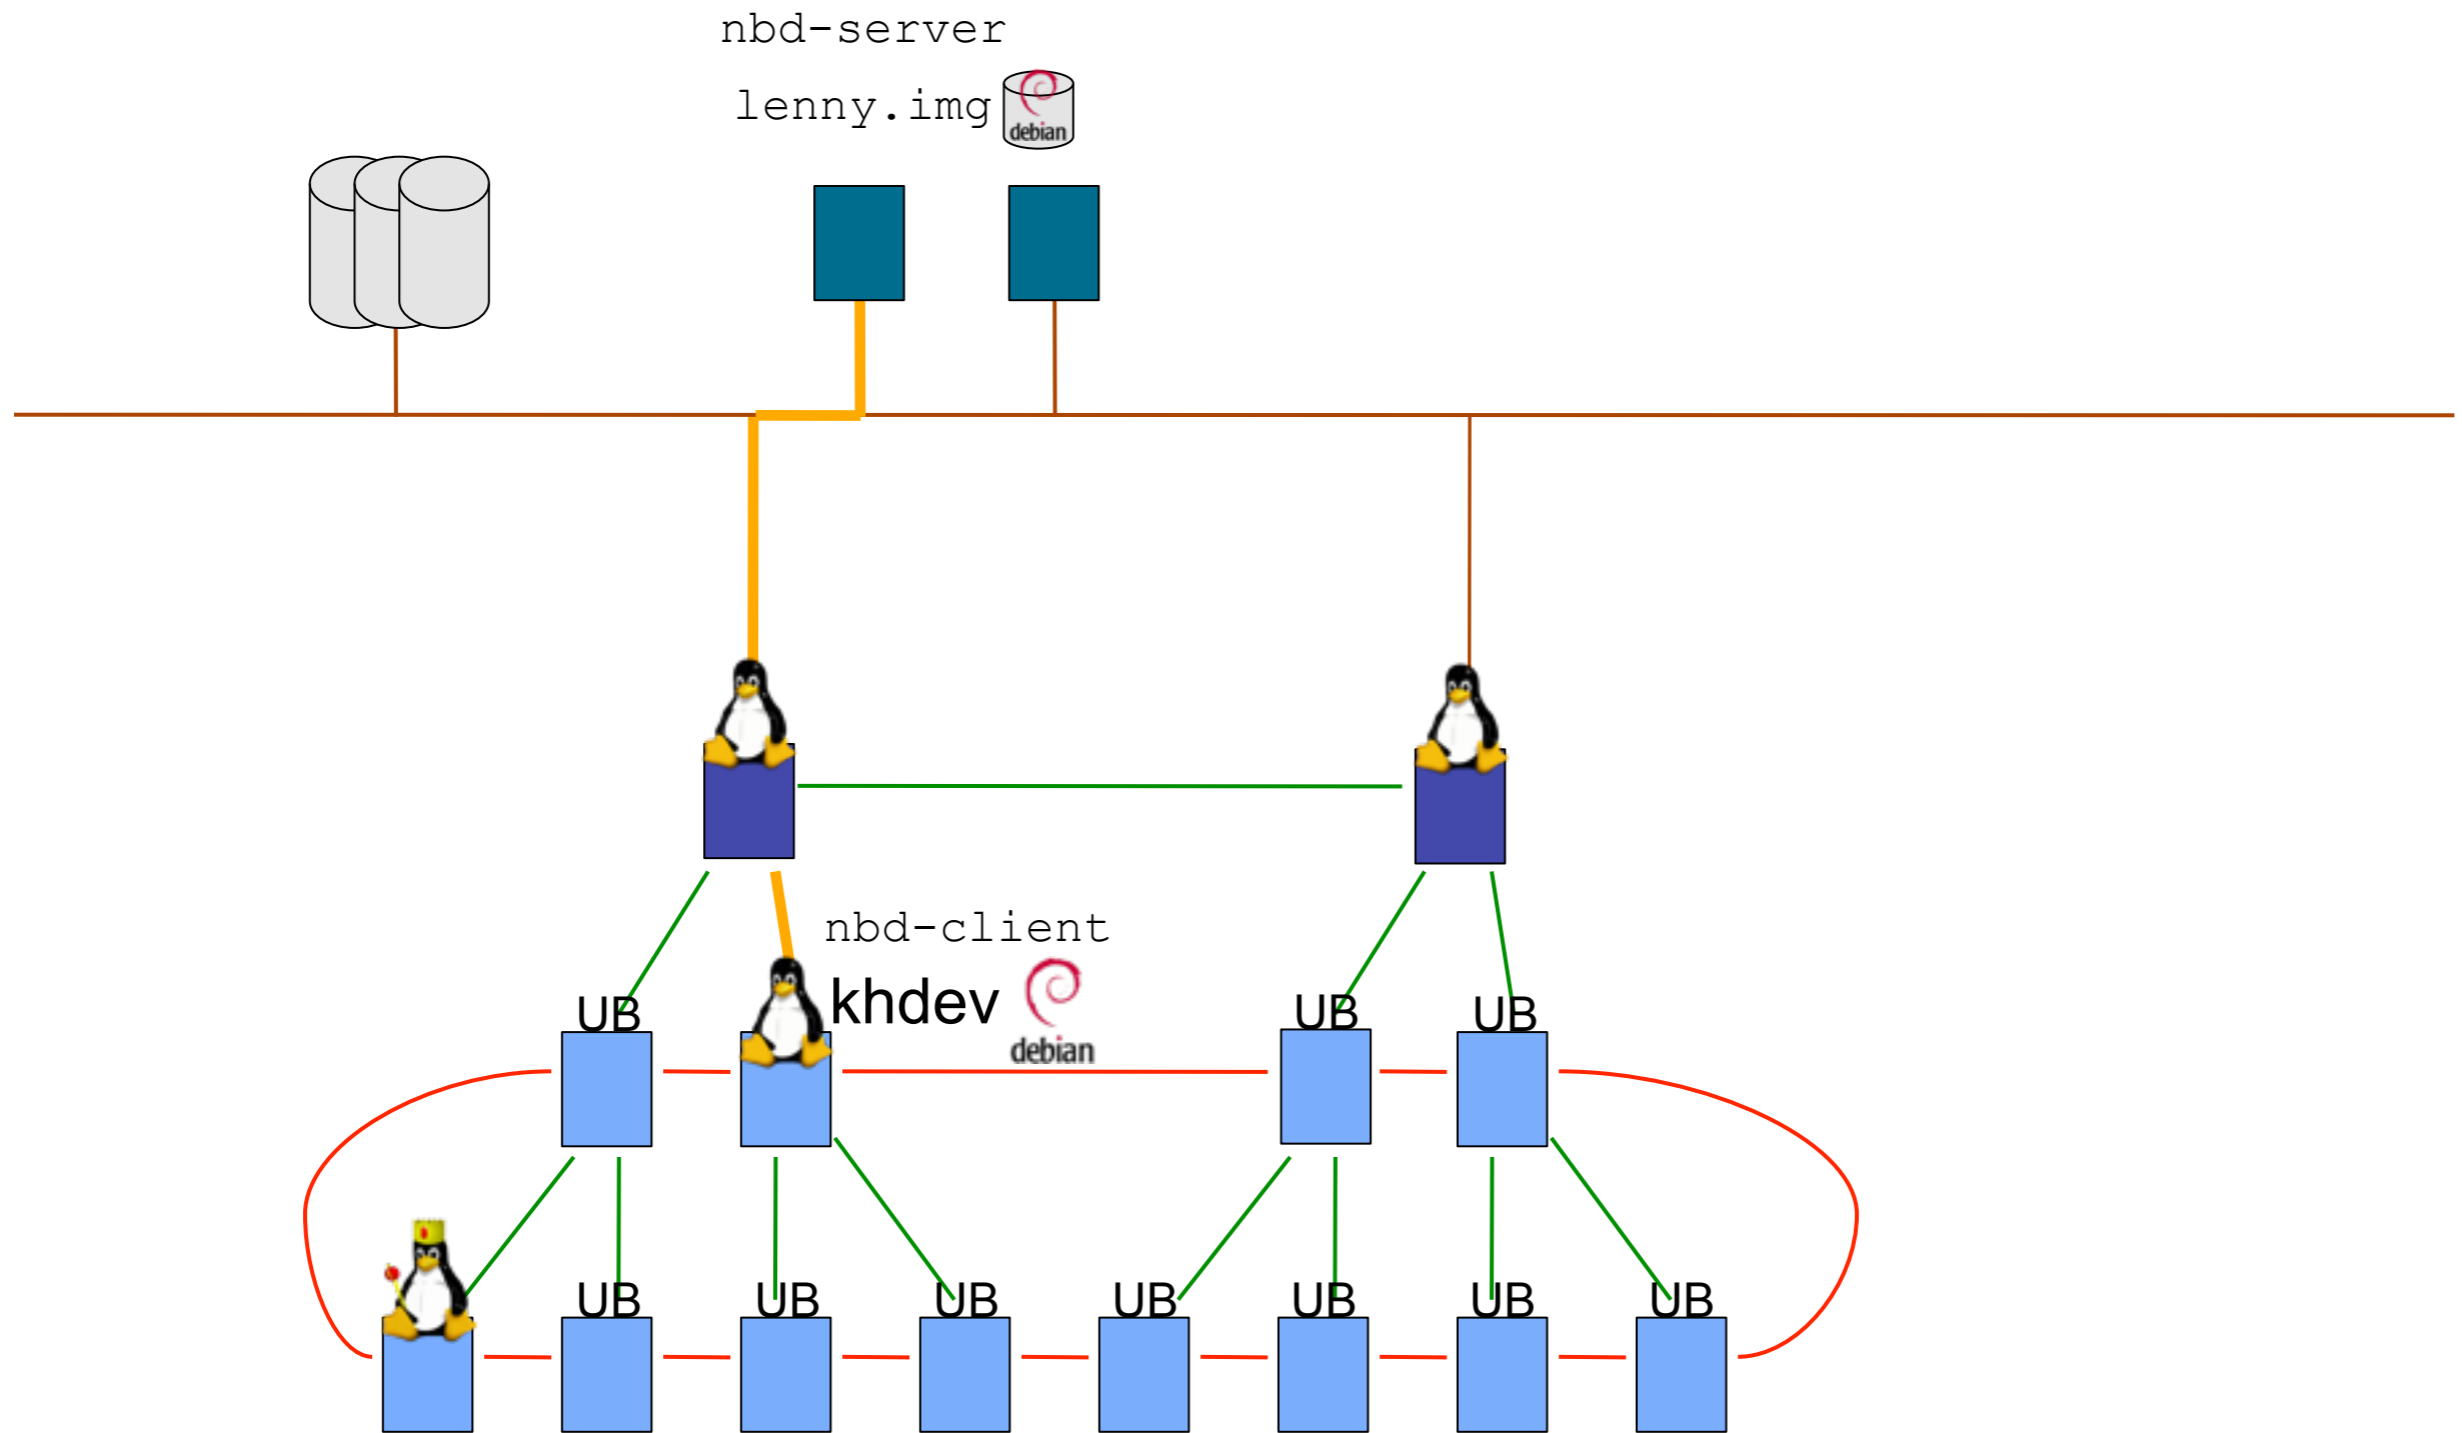
MPI dependencies on Overlay Semantics get exposed



```
foxkill <nodes>
```

Contribution 2: Development Environment on Kittyhawk

Kittyhawk FOX Development Environment




```
khdev  
`khdev' [-h] [-c] [-T] [-k kernel] [-r ramdisk] [-l khctlserver ip] [-K key]  
<root disk image>
```


Status

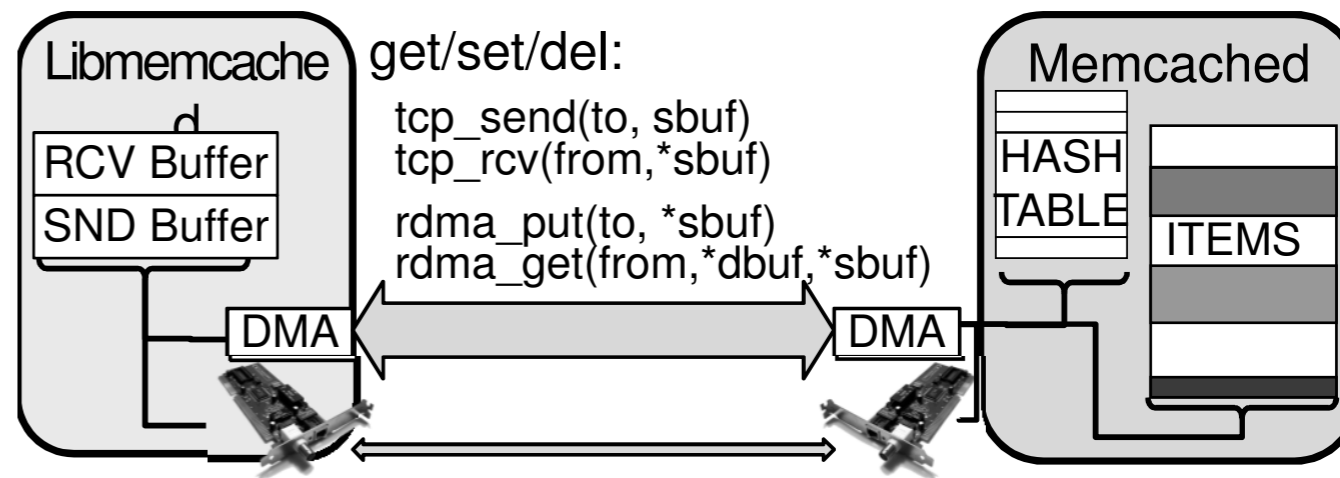
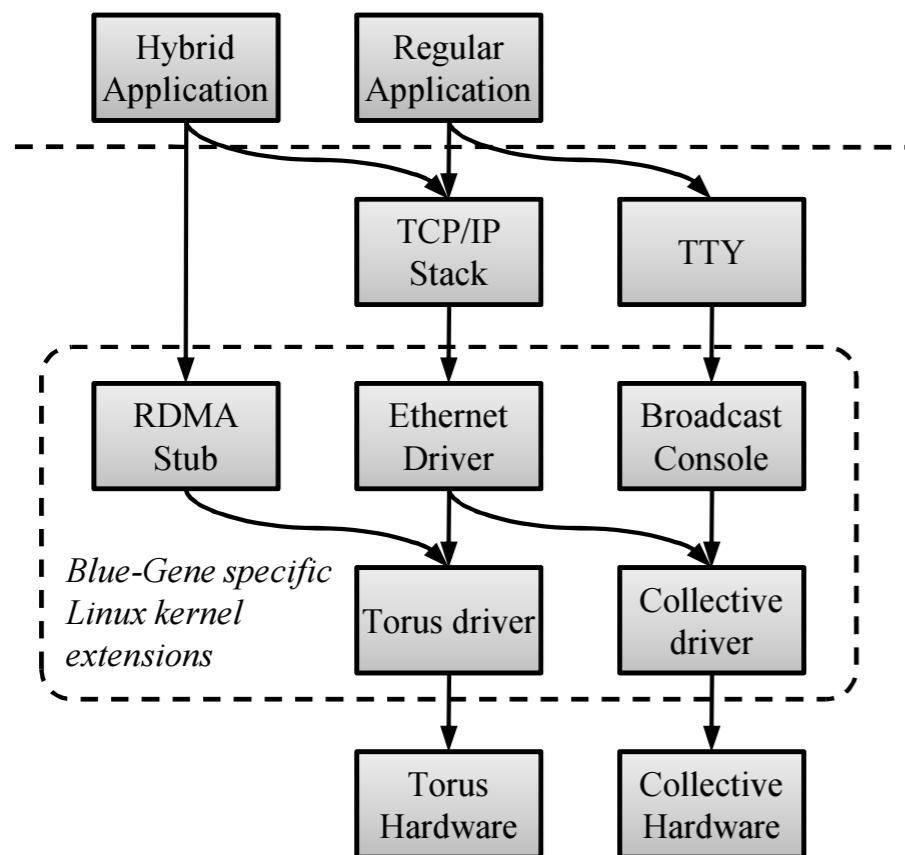
- Improving serial console reliability
 - Adding flow control
- Complete scripting of KHDev environment
- Port to Intrepid/Challenger

Backup



"khstat"
"appstat"

<http://www.cs.bu.edu/~jappavoo/Resources/FOX/khkilldraft.mov>



<http://kittyhawk.bu.edu/Papers/p03.pdf>

Work done with Jan Stoess
 Who is continuing this line of research