Project Blackbird

CLEMSON UNIVERSITY

Utilizing Condor and HTC to address archiving online courses at Clemson on a weekly basis

Sam Hoover

shoover@clemson.edu

LOCKHEED "SKUNK WORKS" SR-71 BLACKBIRD

Project Blackbird

CLEMSON UNIVERSITY

Blackboard at Clemson

- End of Semester archives of all online courses in Blackboard since implementation in 2004
- 77 GB Oracle 10.2.0.4 DB tied to a 1.2 TB Content system with over 13 million files
- Spring 2010: 4610 active Blackboard courses, 31,372 total courses in Blackboard
- Full system backups once a week, nightly incremental backups of entire system

Condor at Clemson

CLEMSON UNIVERSITY

• Clemson has deployed a Condor pool consisting of Windows Vista machines in the public computer labs and several other groups of machines (Linux, Solaris, etc.). These machines are available to Clemson faculty, students, and staff with high-throughput computing needs. Users can create their own Condor submit machines by downloading the appropriate software, and can even contribute their own idle cycles to the pool.

Condor at Clemson

CLEMSON UNIVERSITY

- The Palmetto Cluster is a dedicated Linux cluster of 1111 nodes. Each node has 8 cores and 12-16 GB of RAM.
- Nodes are sold as “Condos” so that the owner gets a guaranteed slice of time based on the number of nodes that they own each week.
- Clemson Condor users get time on the system if it is not in use by a Condo owner.
- We also share cycles via OSG, as the lowest priority user on the system.

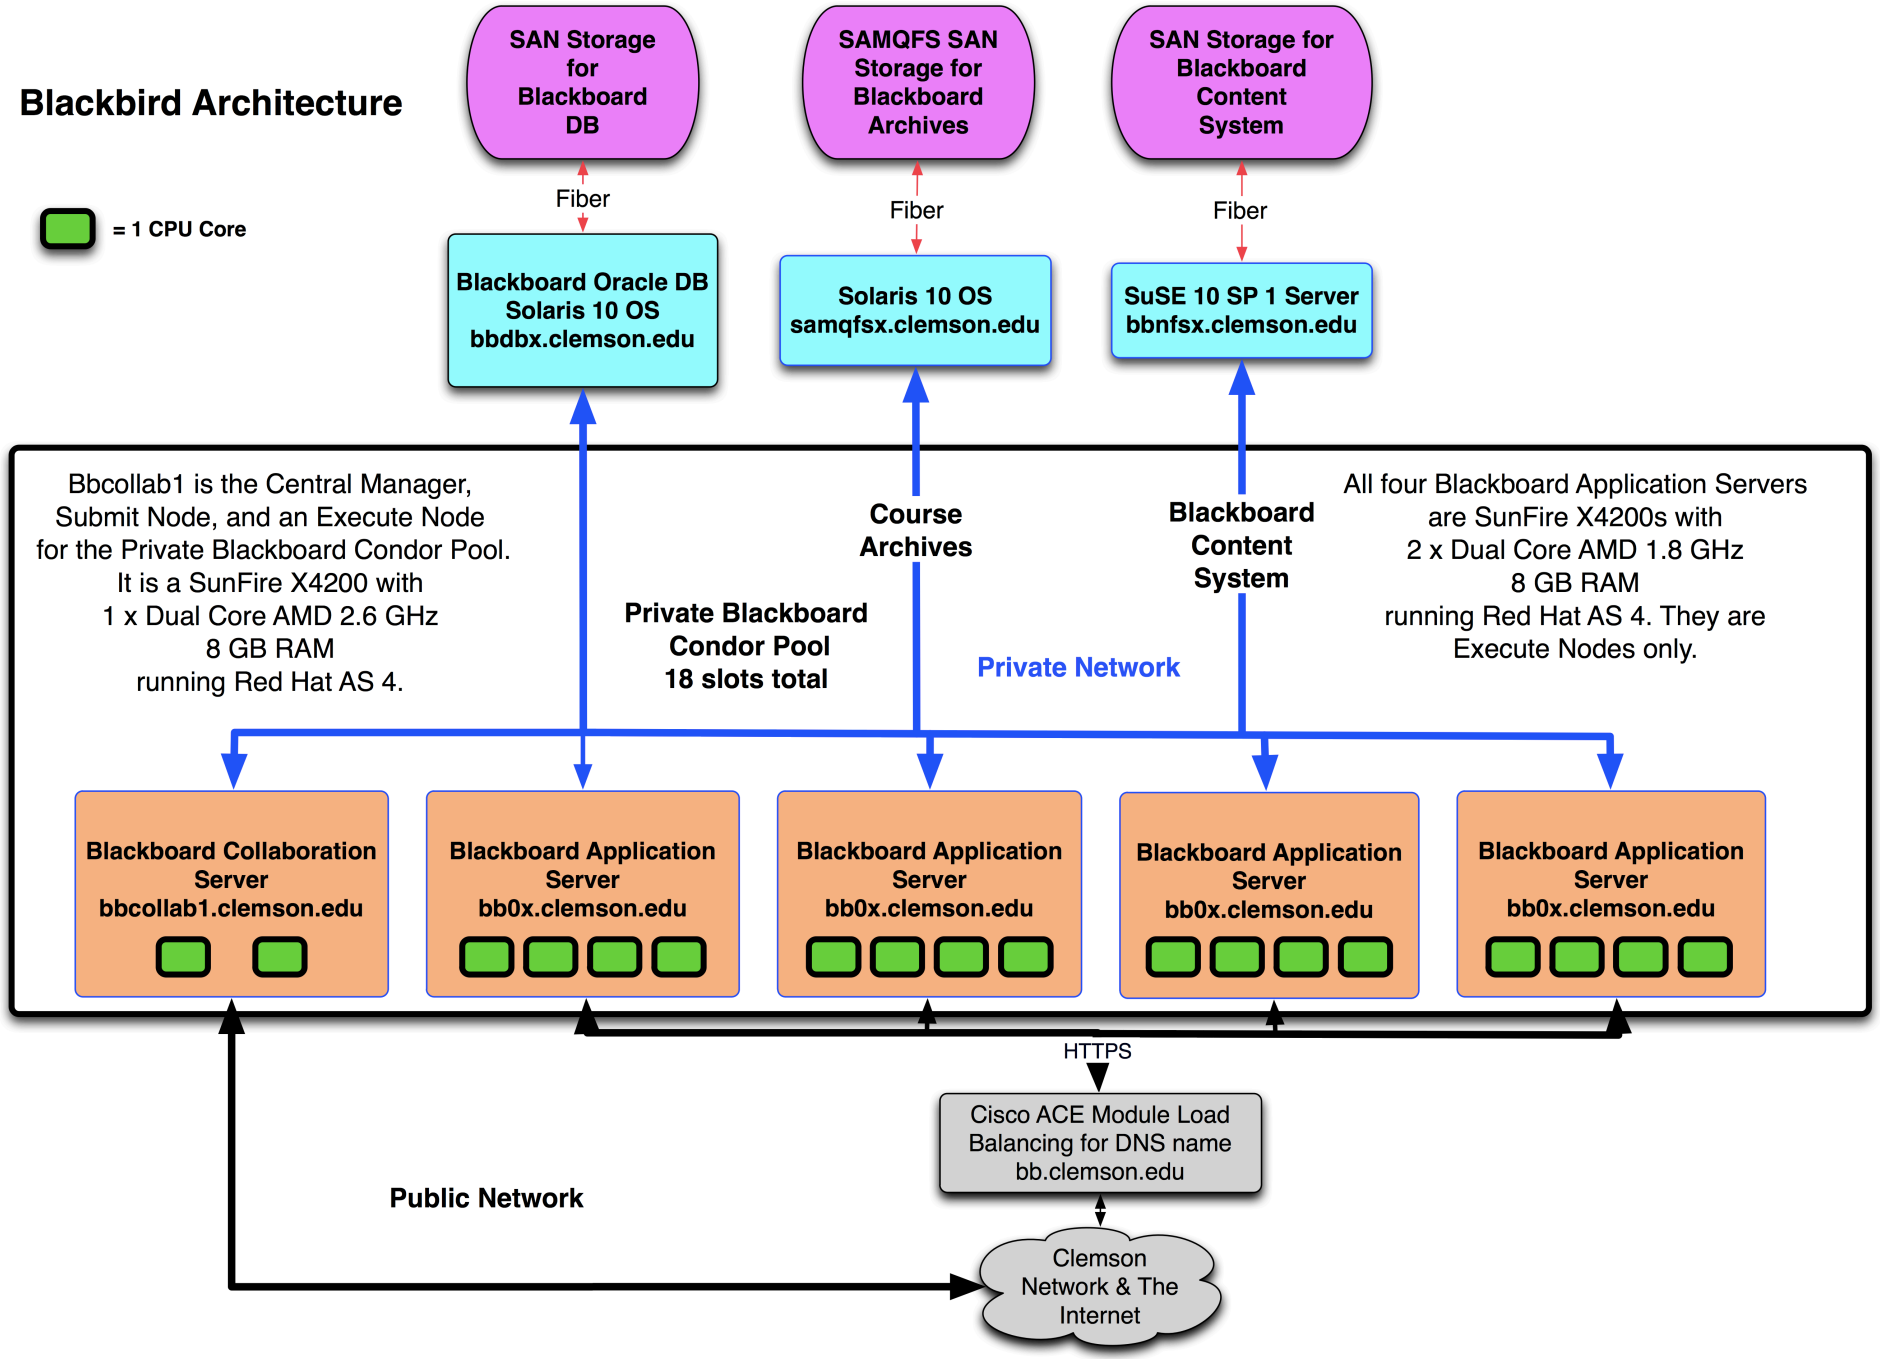
Blackbird Architecture
= 1 CPU Core
SAN Storage for Blackboard DB
SAMQFS SAN Storage for Blackboard Archives
SAN Storage for Blackboard Content System
Fiber
Fiber
Fiber
Blackboard Oracle DB
Solaris 10 OS
bbdbx.clemson.edu
Solaris 10 OS
samqfsx.clemson.edu
SuSE 10 SP 1 Server
bbnfsx.clemson.edu
Bbcollab1 is the Central Manager, Submit Node, and an Execute Node for the Private Blackboard Condor Pool. It is a SunFire X4200 with 1 x Dual Core AMD 2.6 GHz 8 GB RAM running Red Hat AS 4.
Course Archives
Blackboard Content System
All four Blackboard Application Servers are SunFire X4200s with 2 x Dual Core AMD 1.8 GHz 8 GB RAM running Red Hat AS 4. They are Execute Nodes only.
Private Blackboard Condor Pool 18 slots total
Private Network
Blackboard Collaboration Server
bbcollab1.clemson.edu
Blackboard Application Server
bb0x.clemson.edu
Blackboard Application Server
bb0x.clemson.edu
Blackboard Application Server
bb0x.clemson.edu
Blackboard Application Server
bb0x.clemson.edu
HTTPS
Cisco ACE Module Load Balancing for DNS name bb.clemson.edu
Public Network
Clemson Network & The Internet

Blackbird Archive

CLEMSON UNIVERSITY

- Blackboard provides a script for executing batch archives given a list of courses as input.
- Weekly archive process at Clemson began in Fall 2006 after an accidental deletion of many courses.
- Started out splitting the course list into four equal chunks and giving each server ¼ of the total course list. All four servers usually finished within 2 hours of each other, total time for the batch was < 24 hours.
- By Fall 2008, archiving the active courses took 85.5 hours, and the servers finished at widely varying times.

Blackbird Archive

CLEMSON UNIVERSITY

- /usr/local/blackboard/apps/content-exchange/bin/batch_ImportExport.sh
- Archive/Restore: The Archive Course function creates a record of the Course including User interactions. It is most useful for recalling Student performance or interactions at later time. The archive package is saved as a .ZIP file that can be restored to the Blackboard system at another time. In effect, Archive/Restore acts as a backup tool at the individual
course level.

Project Blackbird

CLEMSON UNIVERSITY

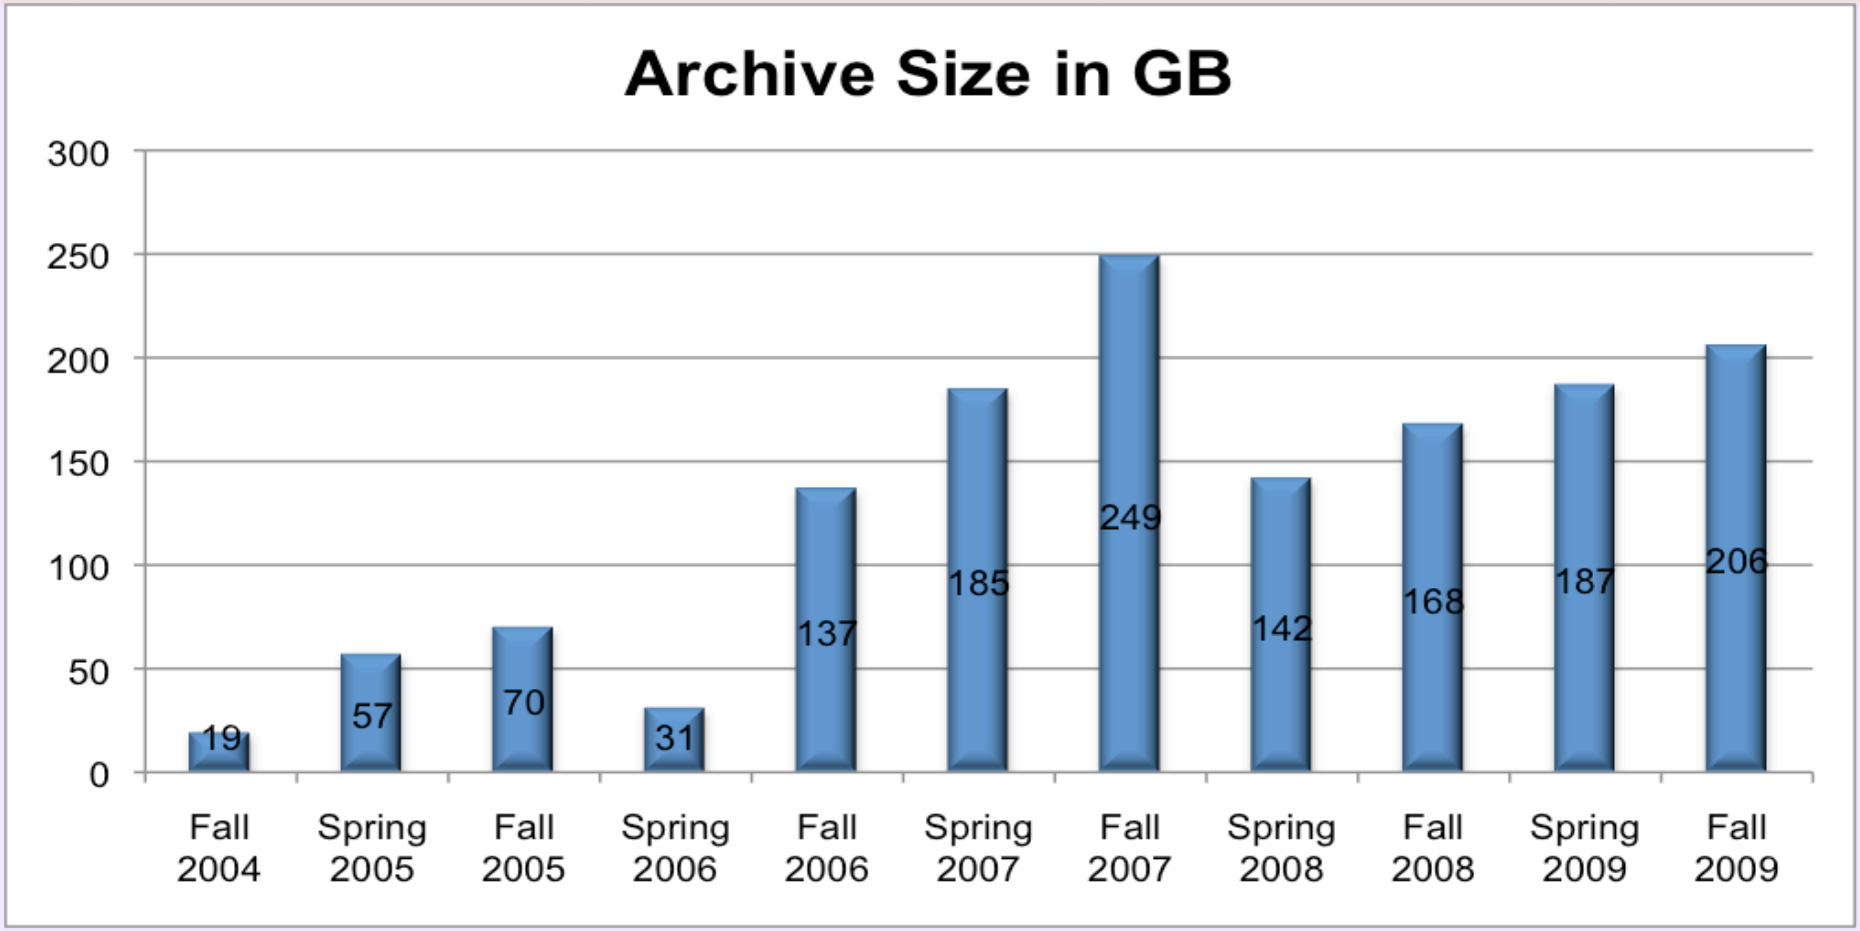

Project Blackbird

CLEMSON UNIVERSITY

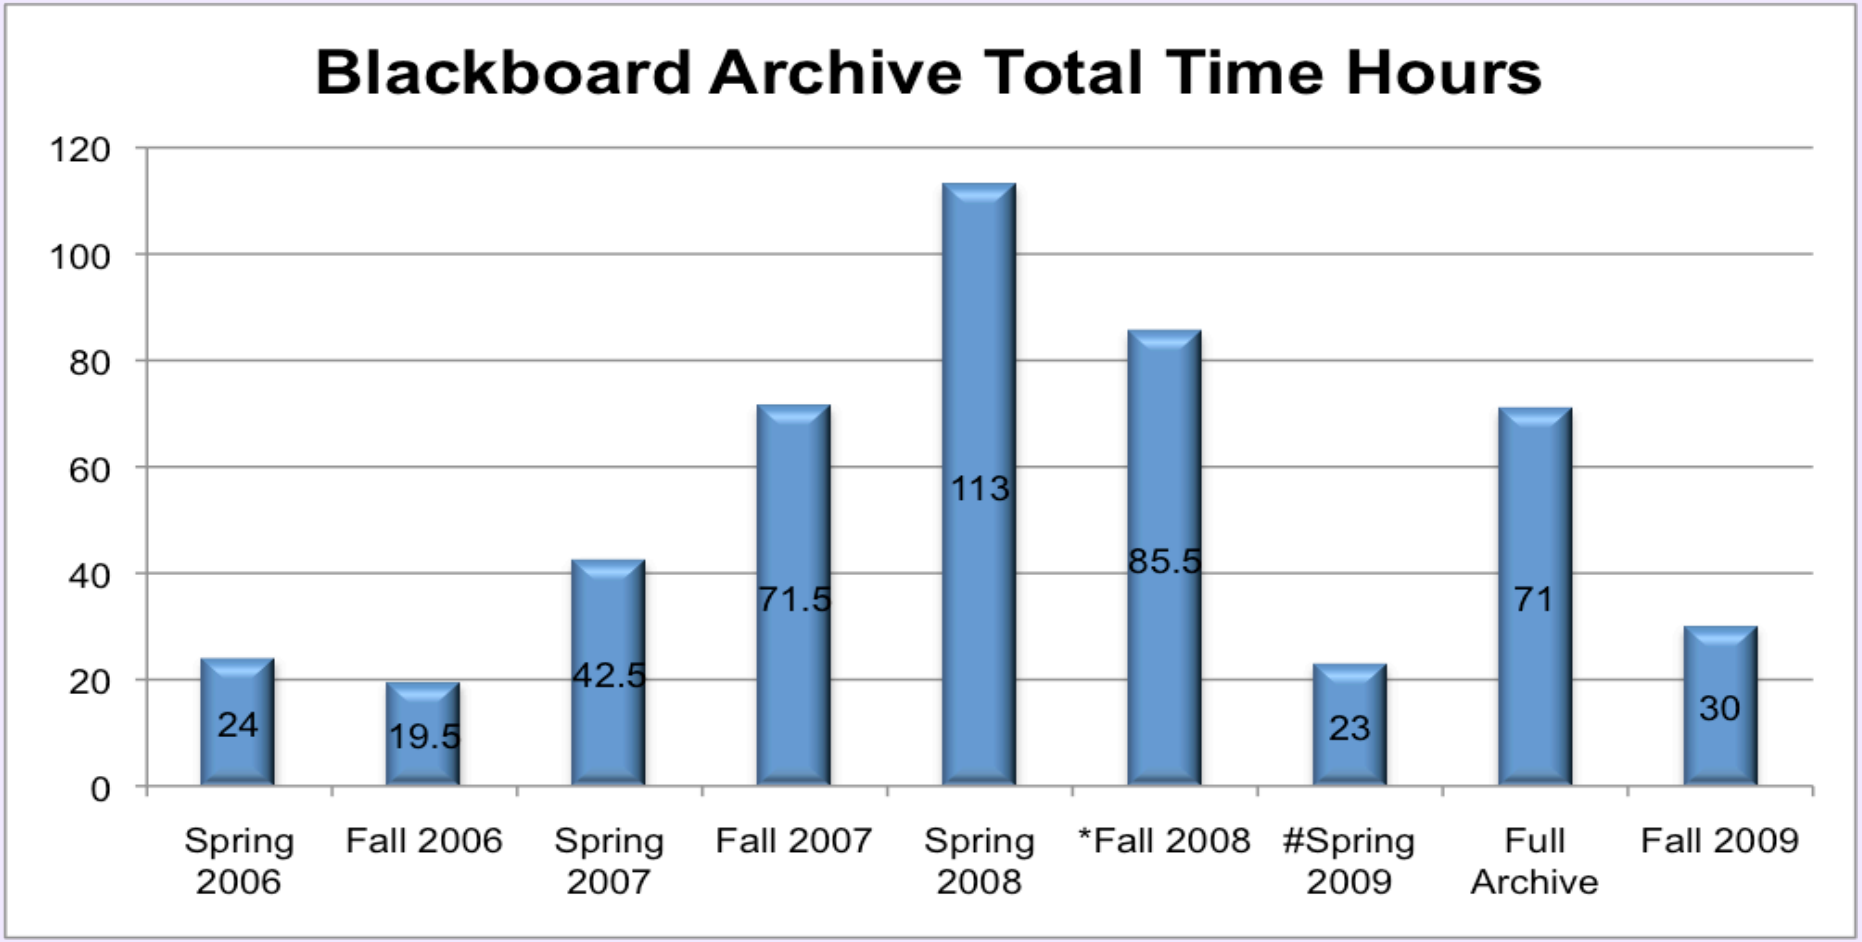

Project Blackbird

CLEMSON UNIVERSITY

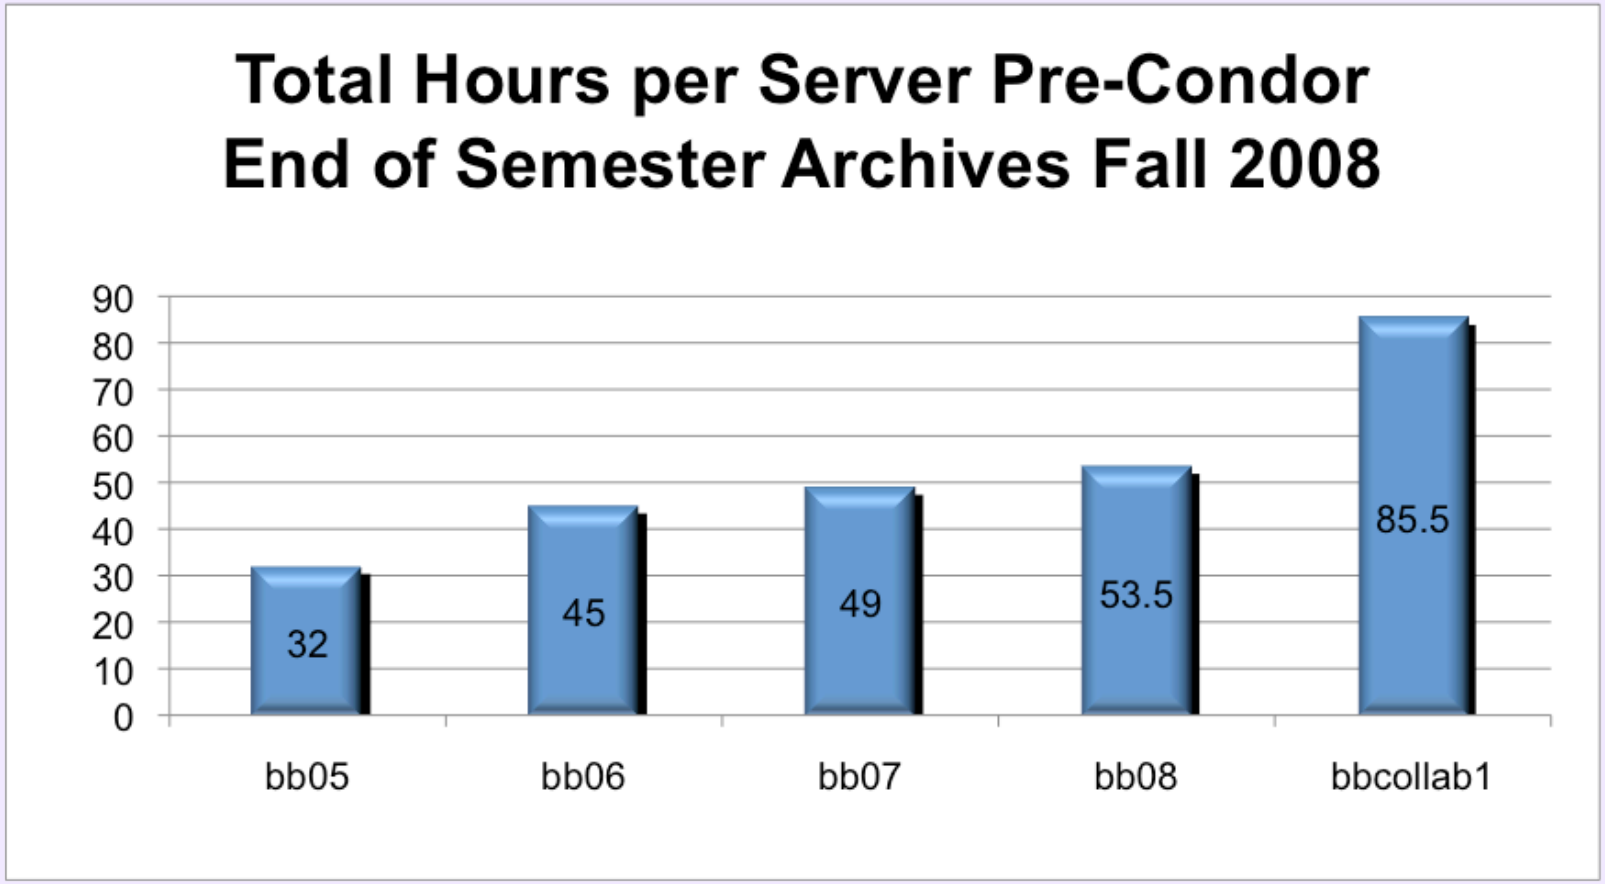

Multiple servers, but 3 cores idle

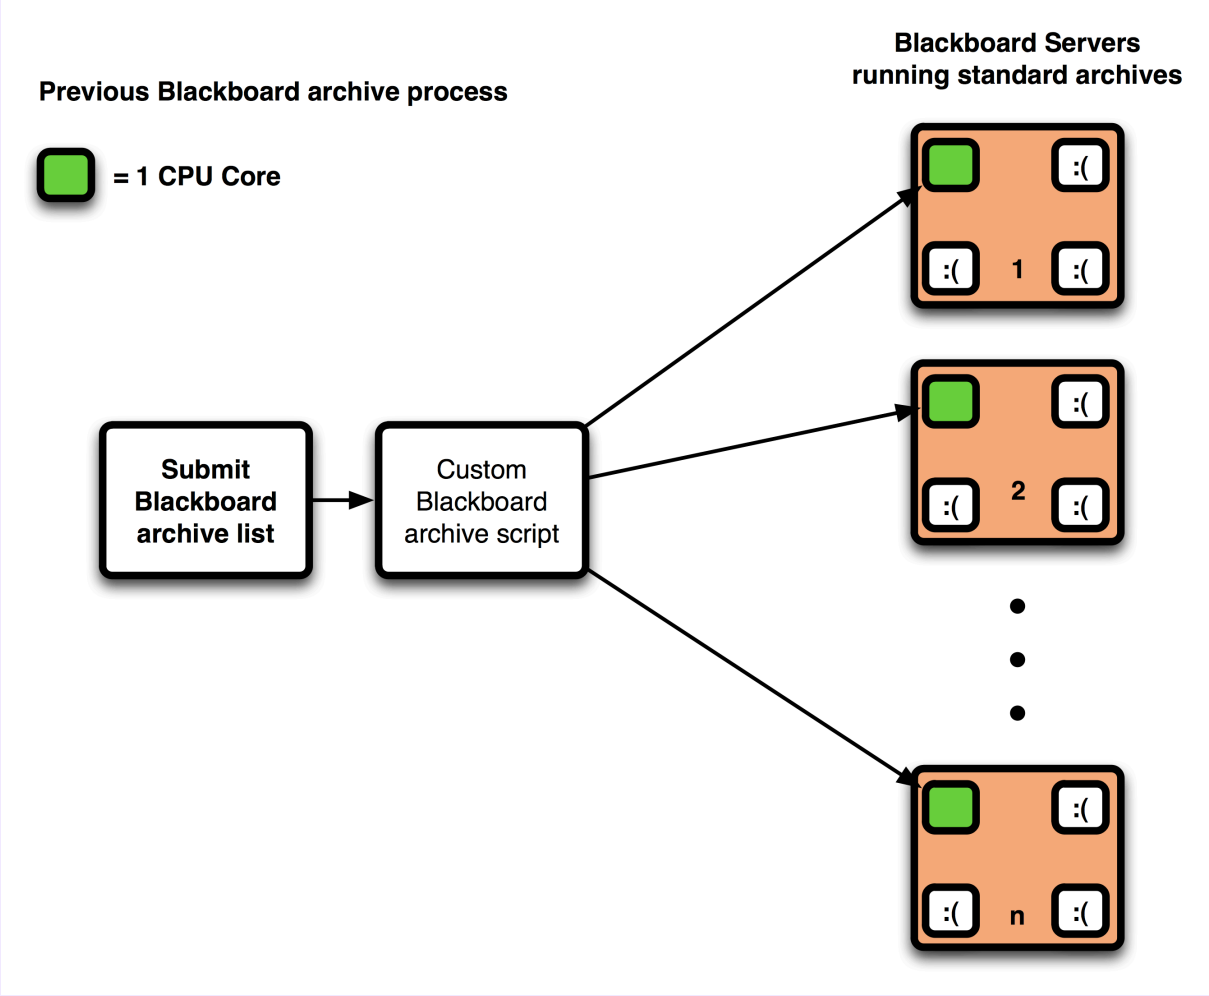

Multiple servers, all cores in use

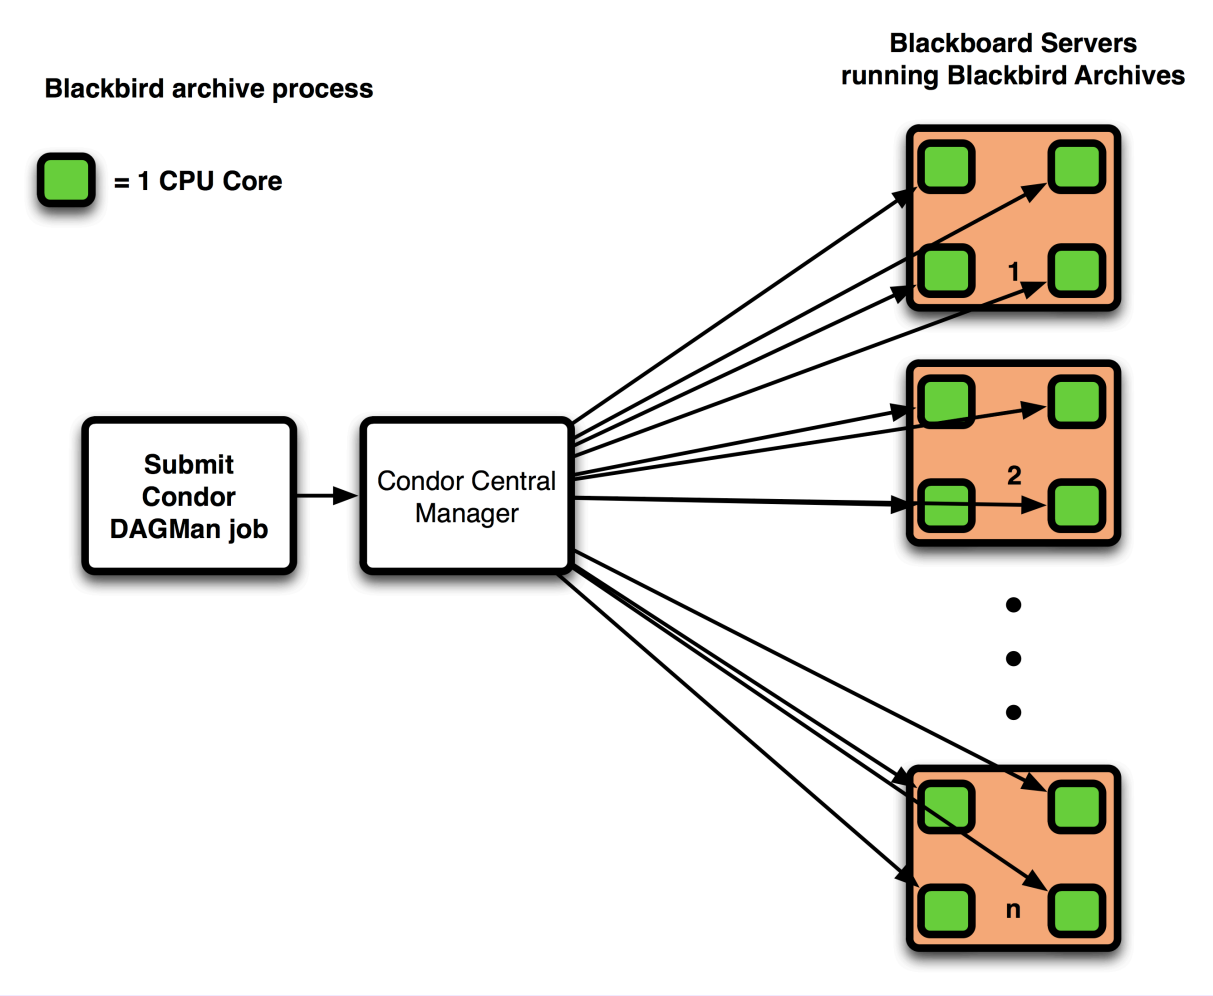

Project Blackbird

CLEMSON UNIVERSITY

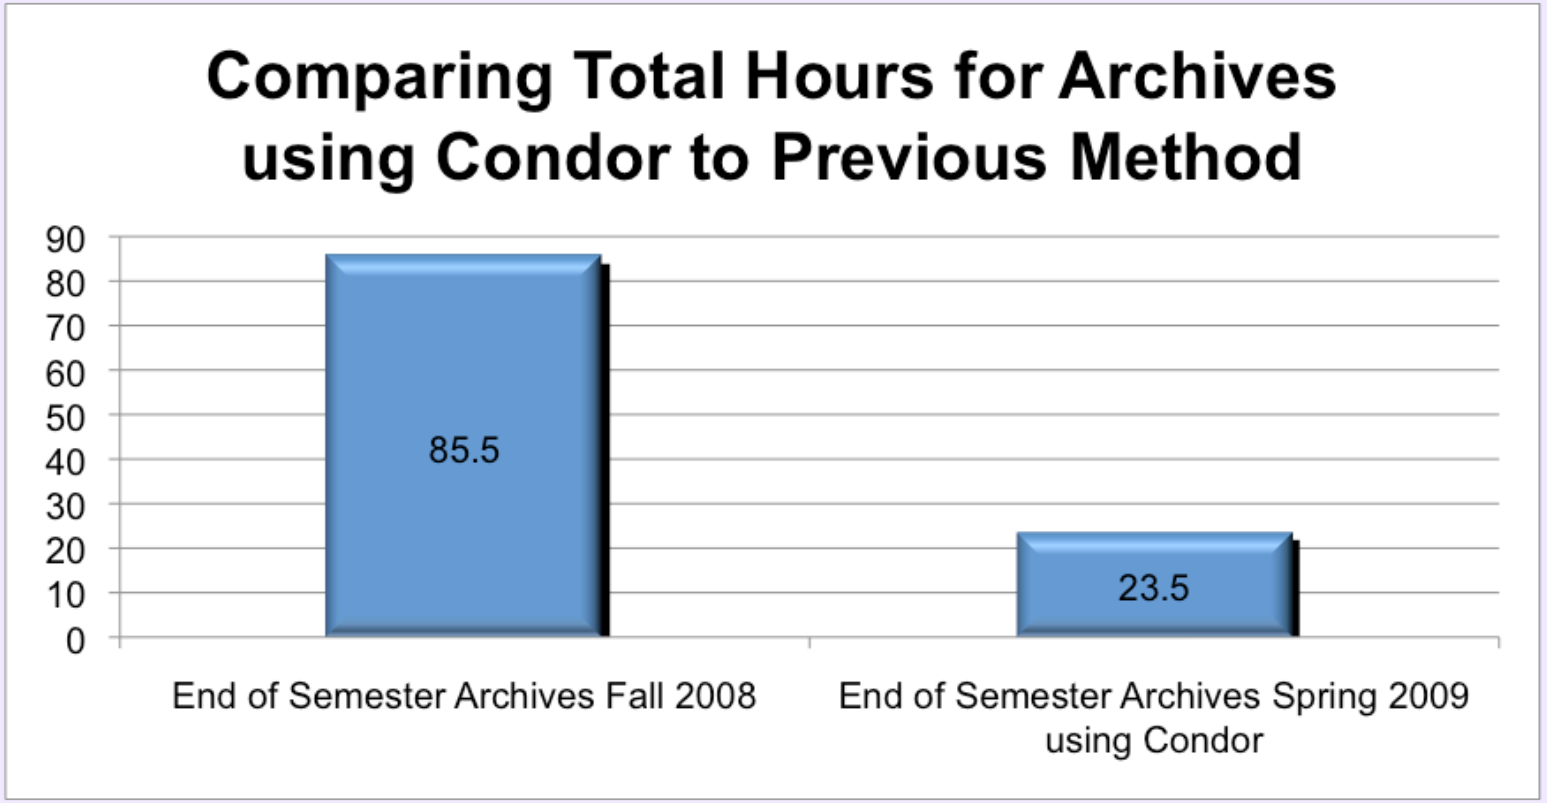

Project Blackbird

CLEMSON UNIVERSITY

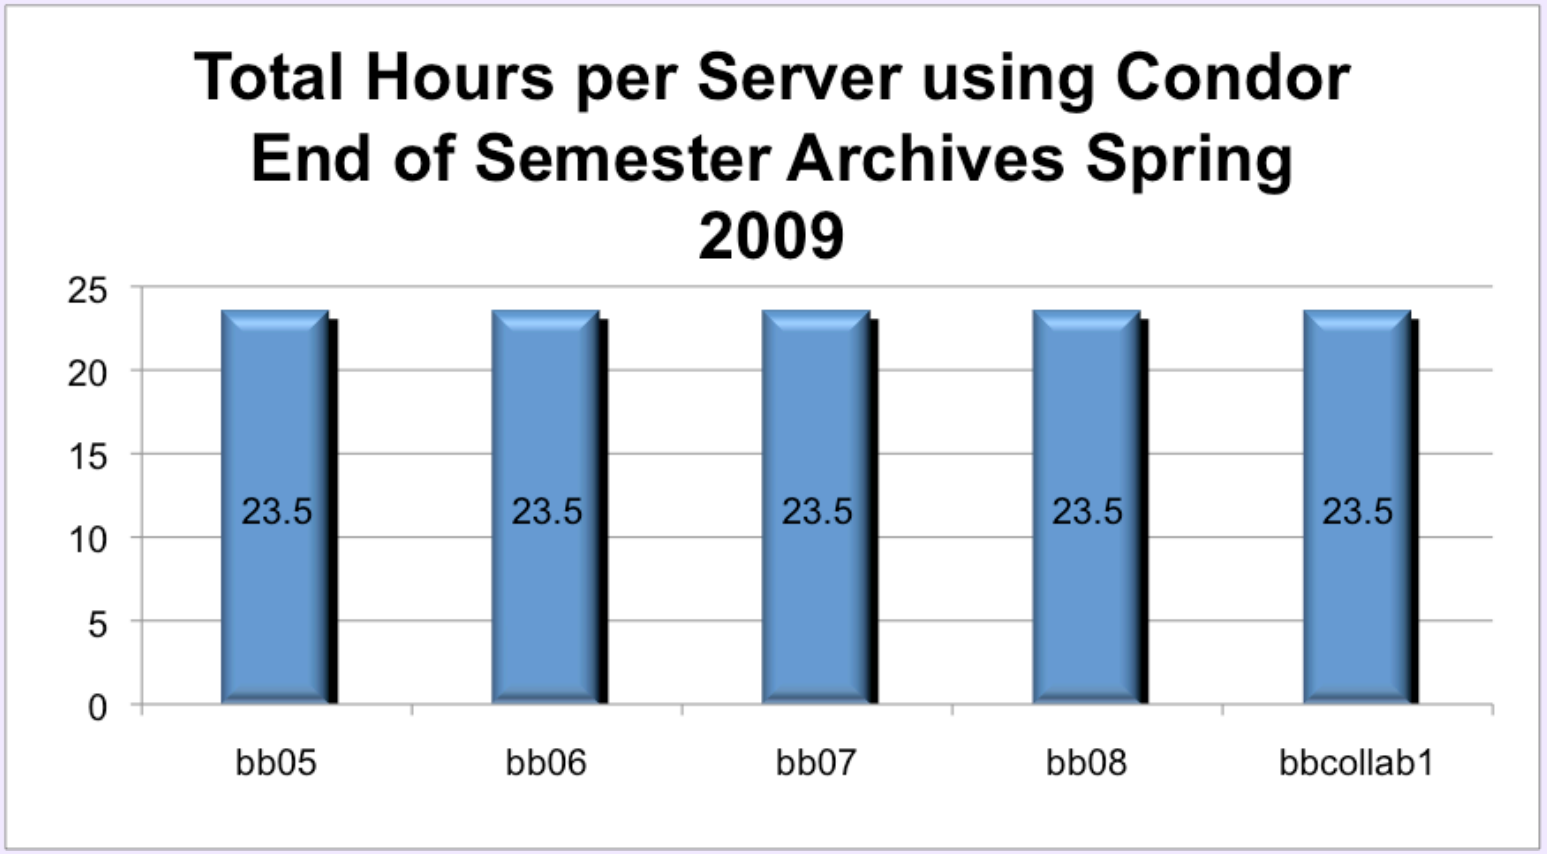

Project Blackbird

CLEMSON UNIVERSITY

Steps in the weekly archive process

- Determine what to archive (active courses, orgs)
- Build a course list
- Create Blackbird submit files
- Submit DAGMan job to Condor
- Monitor Condor queue
- Receive email notification when all courses have been archived

Project Blackbird

CLEMSON UNIVERSITY

What did it take to implement?

- Have one or more multi-core machines
- Choose one machine as your Central Manager
- Install and configure Condor on each machine
- Automate course list creation (Query DB or Directory)
- Automate Condor submit files and Condor DAGMan file creation
- Automate the whole thing with cron
- Check log files for errors upon archive completion

CLEMSON UNIVERSITY

Custom Condor Configuration

```
DAGMAN_MAX_JOBS_IDLE = 25
DAGMAN_MAX_JOBS_SUBMITTED = 50
## Force Condor to use Blackboard Private Network
NETWORK_INTERFACE = Private Blackboard Net
```

DAGMan example

```
# Filename: /usr/local/CMSIntegration/files/Blackbird20091008.condor.sub
# Generated by condor_submit_dag /usr/local/CMSIntegration/files/Blackbird20091008
universe       = scheduler
executable     = /usr/local/condor/bin/condor_dagman
getenv         = True
output         = /usr/local/CMSIntegration/files/Blackbird20091008.lib.out
error          = /usr/local/CMSIntegration/files/Blackbird20091008.lib.err
log            = /usr/local/CMSIntegration/files/Blackbird20091008.dagman.log
remove_kill_sig = SIGUSR1
# Note: default on_exit_remove expression:
# ( ExitSignal =?= 11 || (ExitCode =!= UNDEFINED && ExitCode >=0 && ExitCode <= 2))
# attempts to ensure that DAGMan is automatically
# requeued by the schedd if it exits abnormally or
# is killed (e.g., during a reboot).
on_exit_remove  = ( ExitSignal =?= 11 || (ExitCode =!= UNDEFINED && ExitCode >=0 && ExitCode <= 2))
copy_to_spool   = False
arguments      = "-f -l . -Debug 3 -Lockfile /usr/local/CMSIntegration/files/Blackbird20091008.lock -AutoRescue 1 -
DoRescueFrom 0 -Dag /usr/local/CMSIntegration/files/Blackbird20091008 -CsdVersion $CondorVersion:' '7.2.4' 'Jun' '15'
'2009' 'BuildID:' '159529' '$"
environment    = _CONDOR_DAGMAN_LOG=/usr/local/CMSIntegration/files/
Blackbird20091008.dagman.out;_CONDOR_MAX_DAGMAN_LOG=0
notification    = Complete
queue
```

Condor Submit example

```
universe = vanilla
requirements = (OpSys=="LINUX") && (Memory > 100) && ((Arch=="INTEL")
|| (Arch=="X86_64"))
executable = /usr/local/bin/condorSubmitArchive.pl
arguments = shoover-S0000BKBRD_401001,/san/weeklyArchives/20091008/
getenv = True
log = /usr/local/logs/bbCondorLogs/archive20091008.log
notification = Error
notify_user = DCIT2803_BB_ON_CALL-L@clemson.edu
transfer_executable = False
when_to_transfer_output = ON_EXIT
queue 1
```

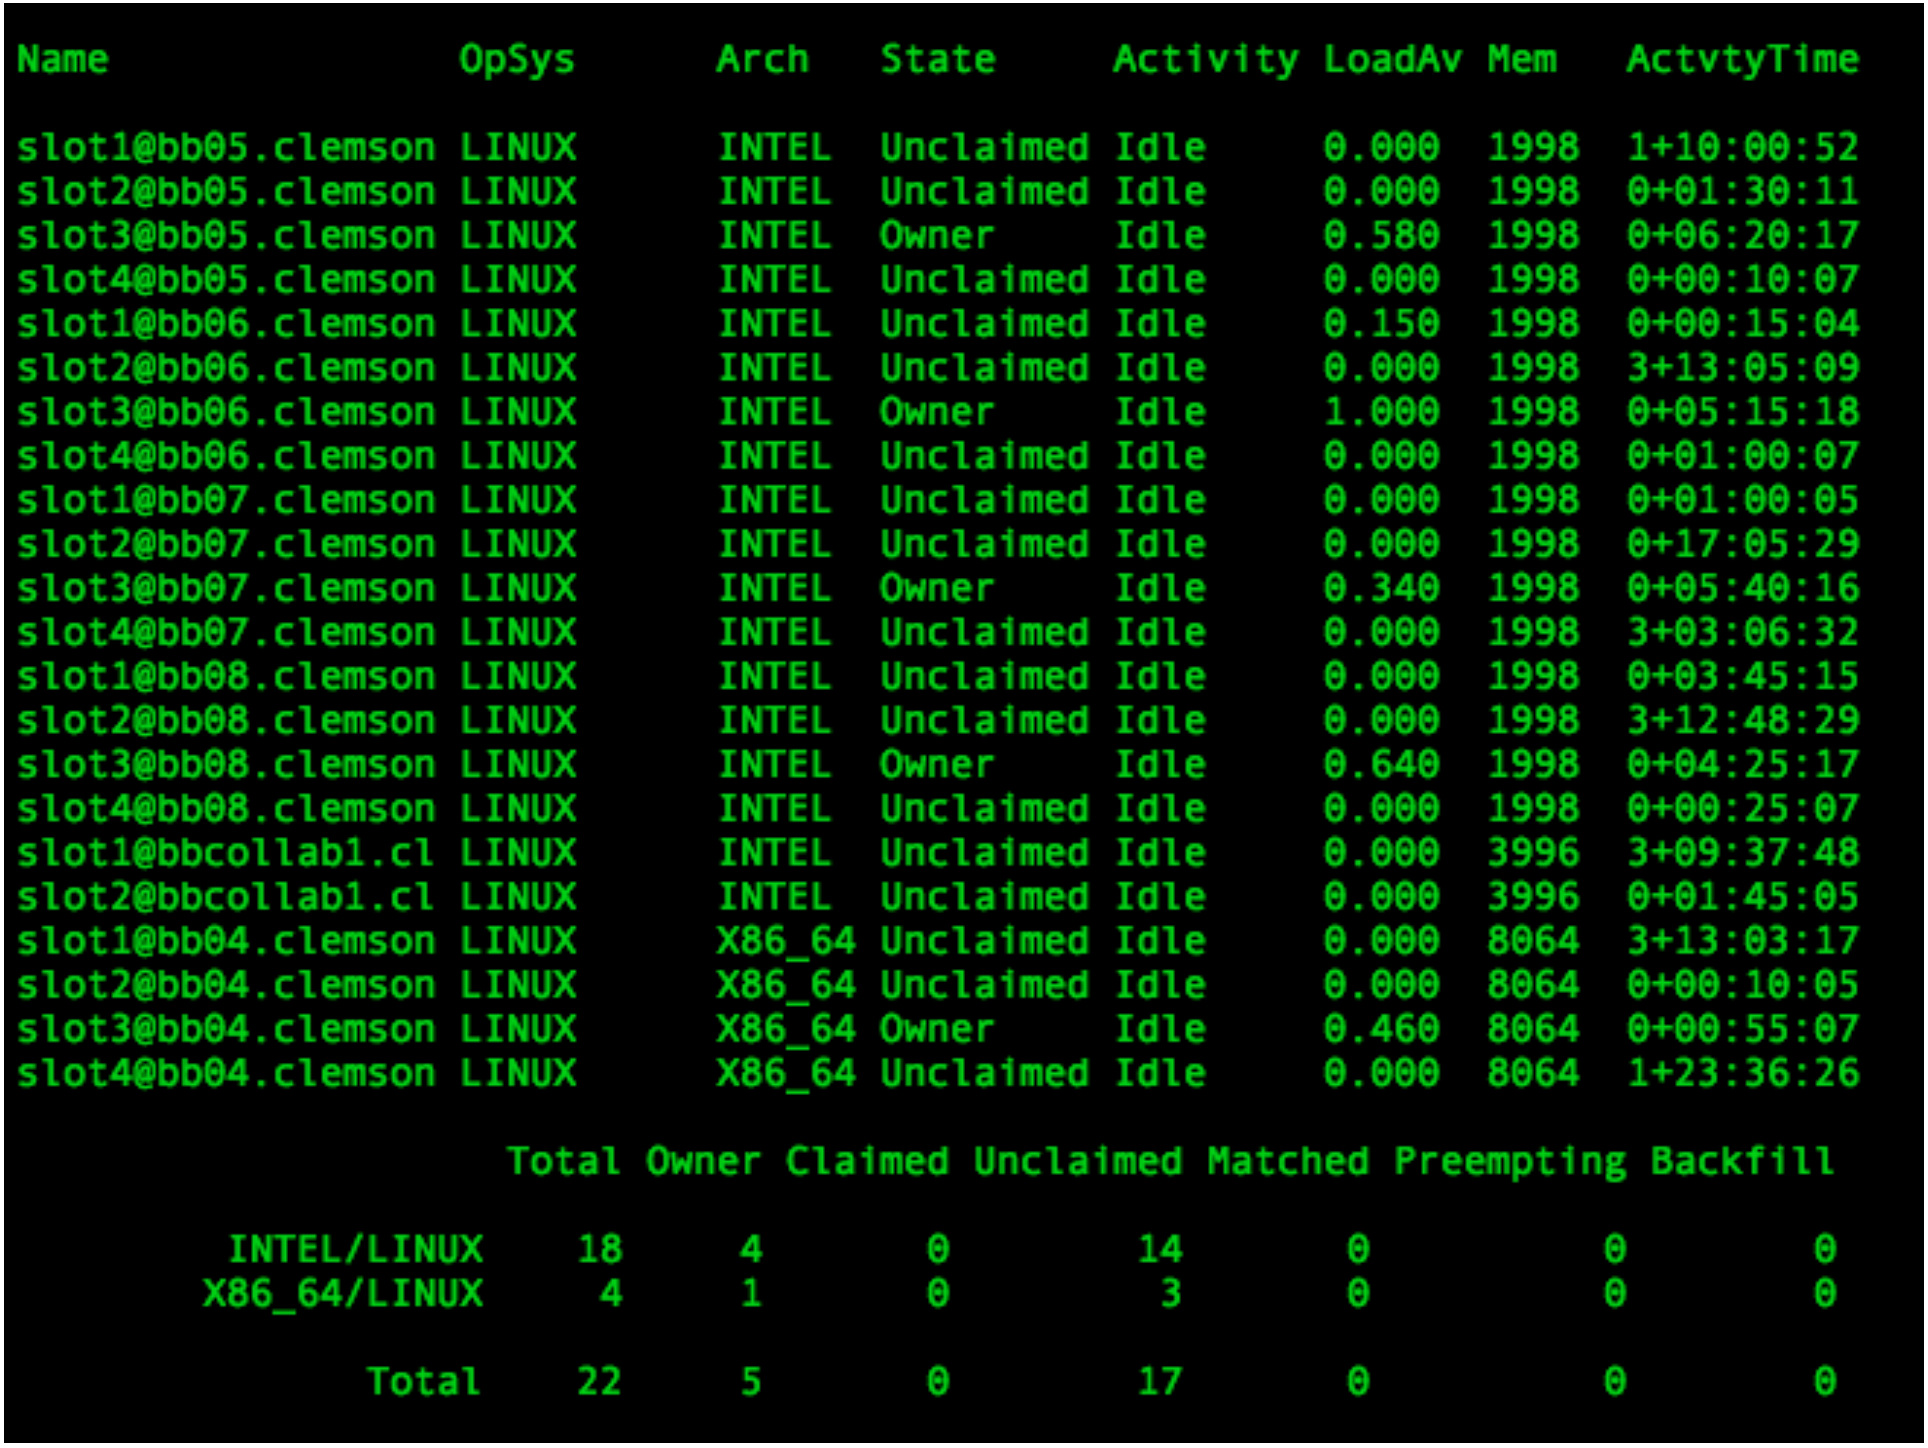

```
Name                OpSys      Arch   State     Activity LoadAv Mem   ActvtyTime

slot1@bb05.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  1+10:00:52
slot2@bb05.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  0+01:30:11
slot3@bb05.clemson LINUX      INTEL  Owner     Idle     0.580  1998  0+06:20:17
slot4@bb05.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  0+00:10:07
slot1@bb06.clemson LINUX      INTEL  Unclaimed Idle     0.150  1998  0+00:15:04
slot2@bb06.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  3+13:05:09
slot3@bb06.clemson LINUX      INTEL  Owner     Idle     1.000  1998  0+05:15:18
slot4@bb06.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  0+01:00:07
slot1@bb07.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  0+01:00:05
slot2@bb07.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  0+17:05:29
slot3@bb07.clemson LINUX      INTEL  Owner     Idle     0.340  1998  0+05:40:16
slot4@bb07.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  3+03:06:32
slot1@bb08.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  0+03:45:15
slot2@bb08.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  3+12:48:29
slot3@bb08.clemson LINUX      INTEL  Owner     Idle     0.640  1998  0+04:25:17
slot4@bb08.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  0+00:25:07
slot1@bbcollab1.cl LINUX      INTEL  Unclaimed Idle     0.000  3996  3+09:37:48
slot2@bbcollab1.cl LINUX      INTEL  Unclaimed Idle     0.000  3996  0+01:45:05
slot1@bb04.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  3+13:03:17
slot2@bb04.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  0+00:10:05
slot3@bb04.clemson LINUX      X86_64 Owner     Idle     0.460  8064  0+00:55:07
slot4@bb04.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  1+23:36:26

                     Total Owner Claimed Unclaimed Matched Preempting Backfill

        INTEL/LINUX    18     4       0        14       0          0        0
      X86_64/LINUX      4     1       0         3       0          0        0

              Total    22     5       0        17       0          0        0
```

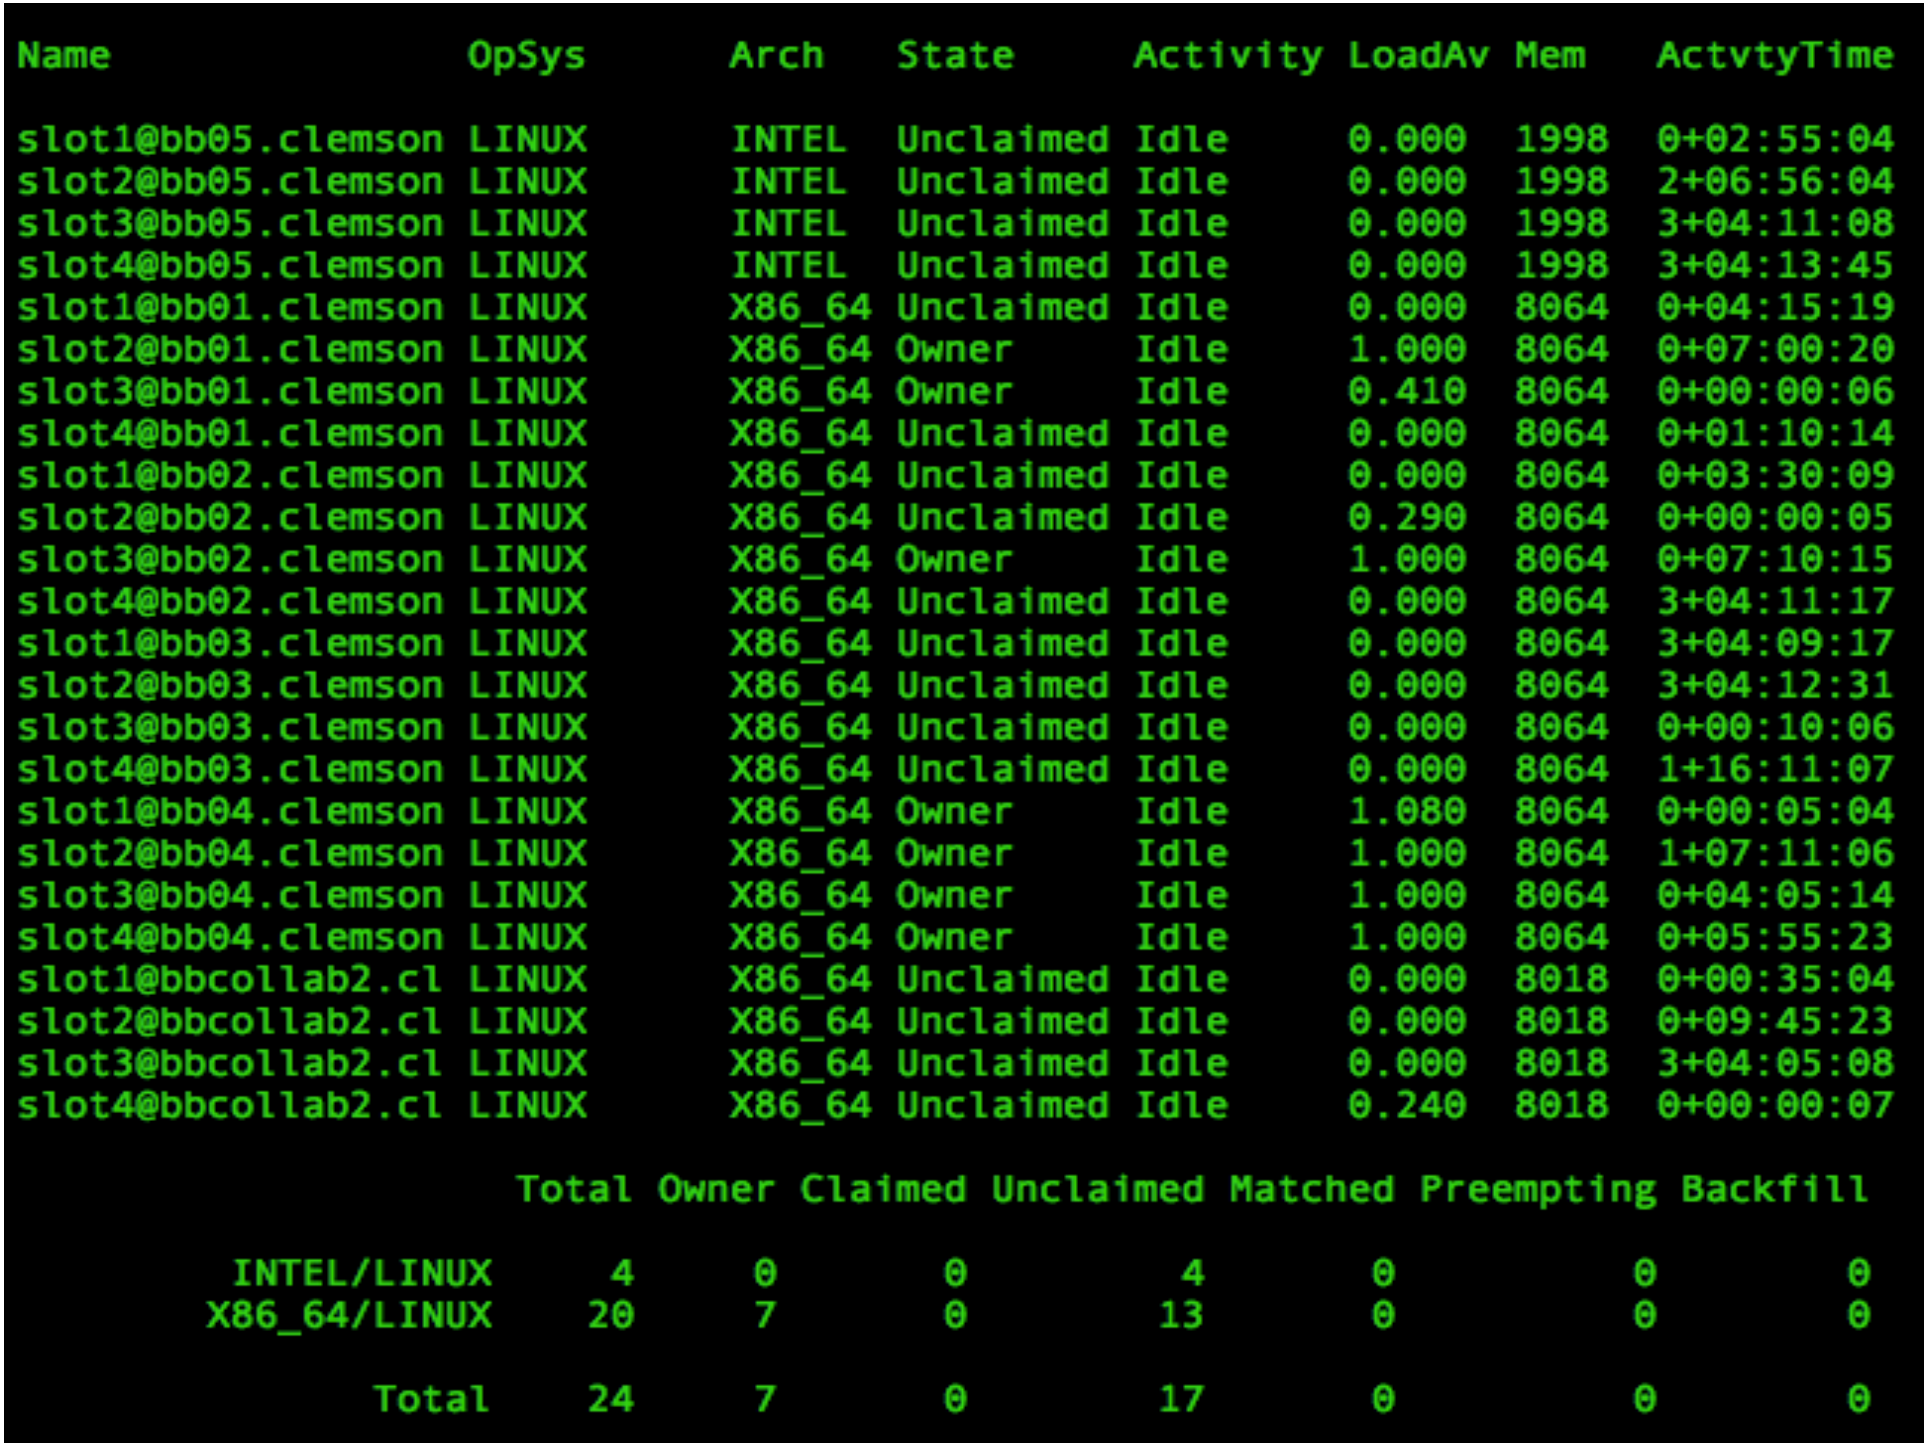

```
Name                OpSys      Arch   State     Activity LoadAv Mem   ActvtyTime

slot1@bb05.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  0+02:55:04
slot2@bb05.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  2+06:56:04
slot3@bb05.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  3+04:11:08
slot4@bb05.clemson LINUX      INTEL  Unclaimed Idle     0.000  1998  3+04:13:45
slot1@bb01.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  0+04:15:19
slot2@bb01.clemson LINUX      X86_64 Owner     Idle     1.000  8064  0+07:00:20
slot3@bb01.clemson LINUX      X86_64 Owner     Idle     0.410  8064  0+00:00:06
slot4@bb01.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  0+01:10:14
slot1@bb02.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  0+03:30:09
slot2@bb02.clemson LINUX      X86_64 Unclaimed Idle     0.290  8064  0+00:00:05
slot3@bb02.clemson LINUX      X86_64 Owner     Idle     1.000  8064  0+07:10:15
slot4@bb02.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  3+04:11:17
slot1@bb03.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  3+04:09:17
slot2@bb03.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  3+04:12:31
slot3@bb03.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  0+00:10:06
slot4@bb03.clemson LINUX      X86_64 Unclaimed Idle     0.000  8064  1+16:11:07
slot1@bb04.clemson LINUX      X86_64 Owner     Idle     1.080  8064  0+00:05:04
slot2@bb04.clemson LINUX      X86_64 Owner     Idle     1.000  8064  1+07:11:06
slot3@bb04.clemson LINUX      X86_64 Owner     Idle     1.000  8064  0+04:05:14
slot4@bb04.clemson LINUX      X86_64 Owner     Idle     1.000  8064  0+05:55:23
slot1@bbcollab2.cl LINUX      X86_64 Unclaimed Idle     0.000  8018  0+00:35:04
slot2@bbcollab2.cl LINUX      X86_64 Unclaimed Idle     0.000  8018  0+09:45:23
slot3@bbcollab2.cl LINUX      X86_64 Unclaimed Idle     0.000  8018  3+04:05:08
slot4@bbcollab2.cl LINUX      X86_64 Unclaimed Idle     0.240  8018  0+00:00:07

                     Total Owner Claimed Unclaimed Matched Preempting Backfill

        INTEL/LINUX      4     0       0         4       0          0        0
      X86_64/LINUX      20     7       0        13       0          0        0

              Total      24     7       0        17       0          0        0
```

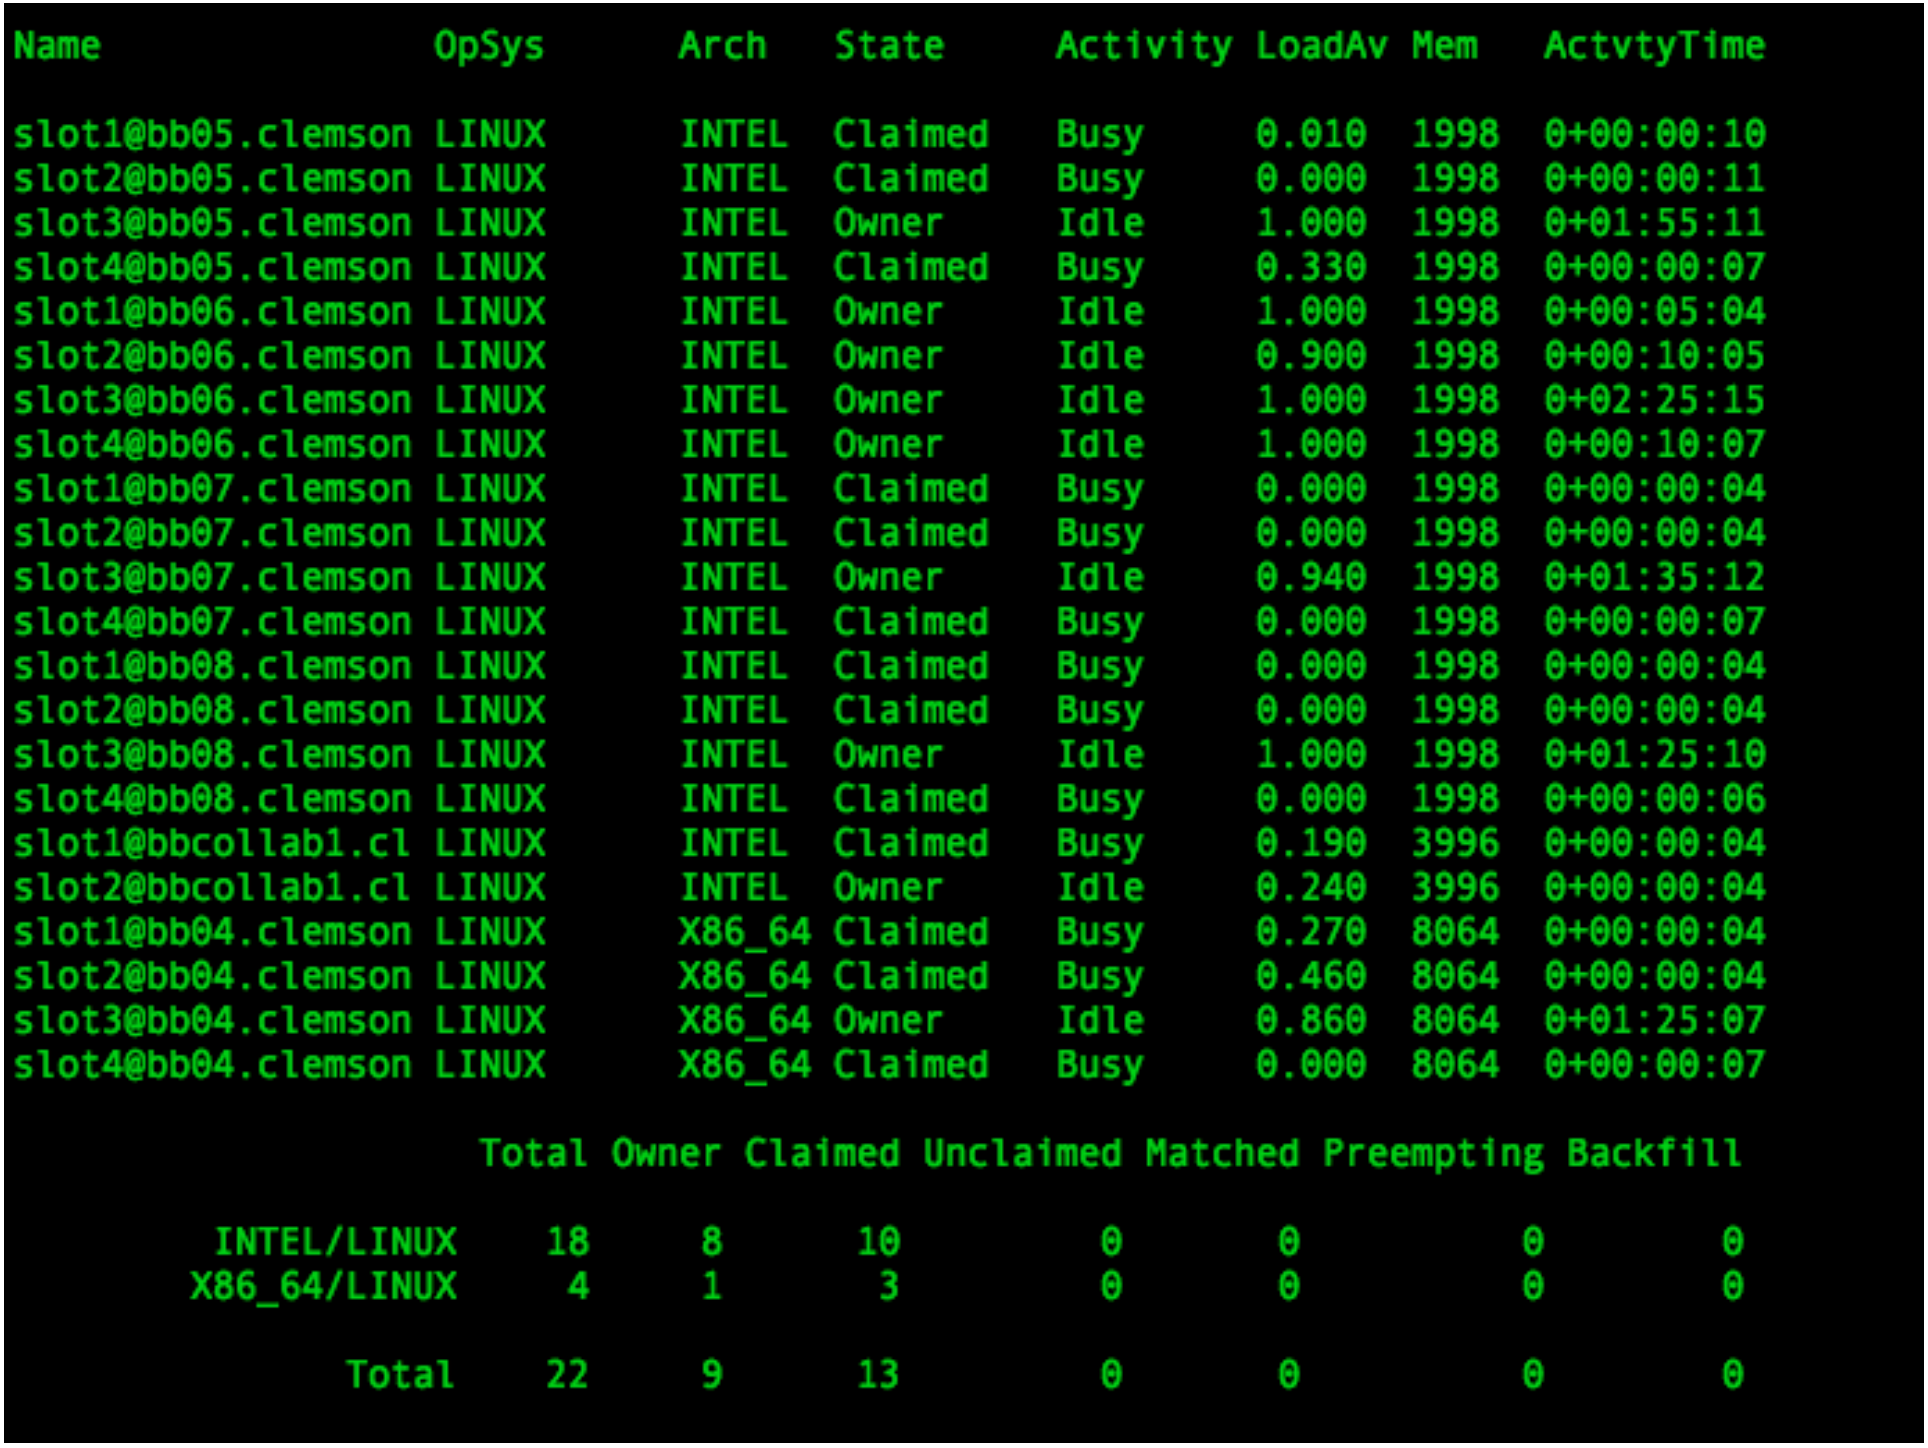

```
Name                OpSys      Arch   State     Activity LoadAv Mem   ActvtyTime

slot1@bb05.clemson LINUX      INTEL  Claimed   Busy     0.010  1998  0+00:00:10
slot2@bb05.clemson LINUX      INTEL  Claimed   Busy     0.000  1998  0+00:00:11
slot3@bb05.clemson LINUX      INTEL  Owner     Idle     1.000  1998  0+01:55:11
slot4@bb05.clemson LINUX      INTEL  Claimed   Busy     0.330  1998  0+00:00:07
slot1@bb06.clemson LINUX      INTEL  Owner     Idle     1.000  1998  0+00:05:04
slot2@bb06.clemson LINUX      INTEL  Owner     Idle     0.900  1998  0+00:10:05
slot3@bb06.clemson LINUX      INTEL  Owner     Idle     1.000  1998  0+02:25:15
slot4@bb06.clemson LINUX      INTEL  Owner     Idle     1.000  1998  0+00:10:07
slot1@bb07.clemson LINUX      INTEL  Claimed   Busy     0.000  1998  0+00:00:04
slot2@bb07.clemson LINUX      INTEL  Claimed   Busy     0.000  1998  0+00:00:04
slot3@bb07.clemson LINUX      INTEL  Owner     Idle     0.940  1998  0+01:35:12
slot4@bb07.clemson LINUX      INTEL  Claimed   Busy     0.000  1998  0+00:00:07
slot1@bb08.clemson LINUX      INTEL  Claimed   Busy     0.000  1998  0+00:00:04
slot2@bb08.clemson LINUX      INTEL  Claimed   Busy     0.000  1998  0+00:00:04
slot3@bb08.clemson LINUX      INTEL  Owner     Idle     1.000  1998  0+01:25:10
slot4@bb08.clemson LINUX      INTEL  Claimed   Busy     0.000  1998  0+00:00:06
slot1@bbcollab1.cl LINUX      INTEL  Claimed   Busy     0.190  3996  0+00:00:04
slot2@bbcollab1.cl LINUX      INTEL  Owner     Idle     0.240  3996  0+00:00:04
slot1@bb04.clemson LINUX      X86_64 Claimed   Busy     0.270  8064  0+00:00:04
slot2@bb04.clemson LINUX      X86_64 Claimed   Busy     0.460  8064  0+00:00:04
slot3@bb04.clemson LINUX      X86_64 Owner     Idle     0.860  8064  0+01:25:07
slot4@bb04.clemson LINUX      X86_64 Claimed   Busy     0.000  8064  0+00:00:07

                     Total Owner Claimed Unclaimed Matched Preempting Backfill

        INTEL/LINUX    18     8      10         0       0          0        0
      X86_64/LINUX      4     1       3         0       0          0        0

              Total    22     9      13         0       0          0        0
```

```
-- Submitter: bbcollab1.clemson.edu : <10.20.4.25:32822> : bbcollab1.clemson.edu
 ID      OWNER            SUBMITTED     RUN_TIME ST PRI SIZE CMD
200633.0   bbuser          10/8  10:26   0+00:01:43 R  0   4.4  condor_dagman
200634.0   bbuser          10/8  10:26   0+00:01:10 R  0   0.0  condorSubmitArchiv
200635.0   bbuser          10/8  10:26   0+00:01:10 R  0   0.0  condorSubmitArchiv
200636.0   bbuser          10/8  10:26   0+00:01:09 R  0   0.0  condorSubmitArchiv
200637.0   bbuser          10/8  10:26   0+00:01:10 R  0   0.0  condorSubmitArchiv
200638.0   bbuser          10/8  10:26   0+00:01:10 R  0   0.0  condorSubmitArchiv
200639.0   bbuser          10/8  10:26   0+00:01:10 R  0   0.0  condorSubmitArchiv
200640.0   bbuser          10/8  10:26   0+00:01:10 R  0   0.0  condorSubmitArchiv
200641.0   bbuser          10/8  10:26   0+00:01:09 R  0   0.0  condorSubmitArchiv
200642.0   bbuser          10/8  10:26   0+00:01:10 R  0   0.0  condorSubmitArchiv
200643.0   bbuser          10/8  10:26   0+00:01:09 R  0   0.0  condorSubmitArchiv
200645.0   bbuser          10/8  10:26   0+00:00:31 R  0   0.0  condorSubmitArchiv
200648.0   bbuser          10/8  10:26   0+00:00:26 R  0   0.0  condorSubmitArchiv
200649.0   bbuser          10/8  10:26   0+00:00:19 R  0   0.0  condorSubmitArchiv
200650.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200651.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200652.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200653.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200654.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200655.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200656.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200657.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200658.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200659.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200660.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200661.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200662.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200663.0   bbuser          10/8  10:26   0+00:00:00 I  0   0.0  condorSubmitArchiv
200664.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200665.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200666.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200667.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200668.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200669.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200670.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200671.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200672.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200673.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200674.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200675.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200676.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200677.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv
200678.0   bbuser          10/8  10:27   0+00:00:00 I  0   0.0  condorSubmitArchiv

43 jobs; 29 idle, 14 running, 0 held
```

Blackbird Benefits

CLEMSON UNIVERSITY

- Reduced total archive time from > 85 hrs to < 24 hrs
- Job scheduling – all servers finish at the same time
- Zero impact to Blackboard Performance
- Automatic suspension/resumption of archives if Load reaches threshold on any core
- Email notification upon completion of all archives
- Load balancing – archive jobs are distributed as cores become available
- Takes advantage of all available CPU cores instead of just one core per server
- Use ClassAds to specify architecture and memory requirements for large archive jobs

Project Blackbird

CLEMSON UNIVERSITY

Recent Updates

- 64 Bit Red Hat 5.4 OS and JVM 1.6
- Maximum (affordable) RAM per machine – 32 GB
- Web page to view queue and status

What's next?

- Add out of warranty machines to the Blackboard Condor Pool (keep users off of them)
- Monitoring of queue
- Automate installation and configuration

Project Blackbird

CLEMSON UNIVERSITY

Questions?

Sam Hoover
shoover@clemson.edu