Creating a Virtual, Secure Blade Cluster Using Red Hat Enterprise Linux 5

T. Lynn Riggs, Ph.D.
Center of Economic Studies
US Bureau of the Census

This work is unofficial and thus has not undergone the review accorded to official Census Bureau publications. The views expressed in the paper are those of the author and not necessarily those of the U.S. Census Bureau.

USCENSUSBUREAU

Road Map

The RDC Environment

The Transition to Blades

Improvements to the Blades

The RDC Network

The purpose is to provide secure access to confidential Census Bureau and other federal statistical data to authorized researchers on approved projects.

The RDCs are operated as Joint Partnerships between the Census Bureau and leading universities and research institutions.

Security is Paramount

Titles 13 (Census) /26 (IRS) U.S.C. and CIPSEA protect confidentiality

- microdata protected by law

Public perception of improper use of data could …

- reduce response rates
- induce Congress to cut funding or program

Customer Service is Important

- Partner Support is Vital
  - Influential institutions and users
  - Word-of-mouth
- Good User Experience is Important
  - Users' time and money at stake
    - Limited time to spend in the RDC
    - Some travel great distances to use an RDC
  - Important career milestones (dissertation, tenure)

Census Research Data Centers

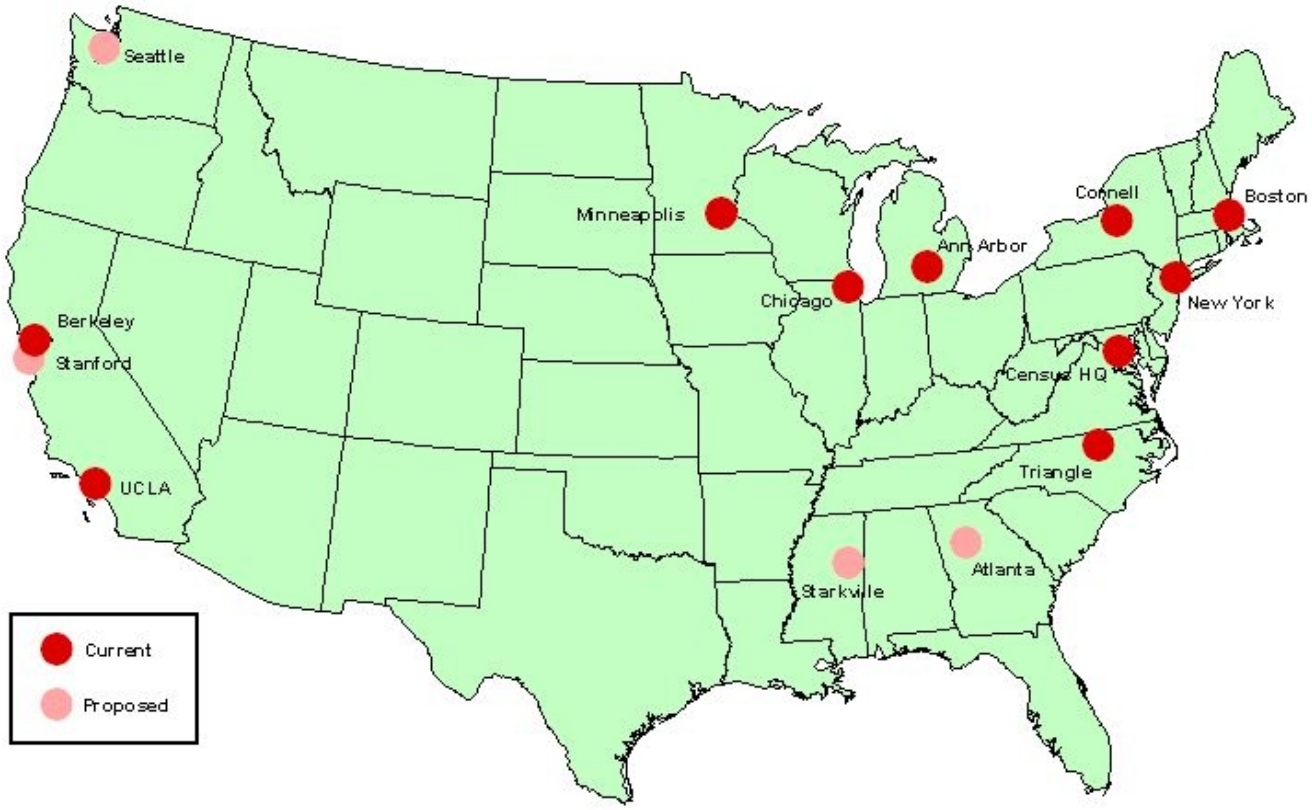

Variable Load Based on RDC Network Activity

135 active projects with > 400 active researchers across the U.S.

24/7 access at each RDC

Computing requirements for each project vary greatly

- Data sets
- Software (and programming skill of user)
- Data manipulation and estimation techniques

Previous RDC Environment

- 15 Different Stand-Alone Servers
  - Dedicated server for each RDC (10 servers)
    - 8 GB of memory
    - ½ TB local storage (local project files)
    - 32-bit
- Central SAN for common files
  - ~ 35TB
  - Read-only data
- Thin client devices connect to the servers using NX

RED HAT :: CHICAGO :: 2009 SUMMIT

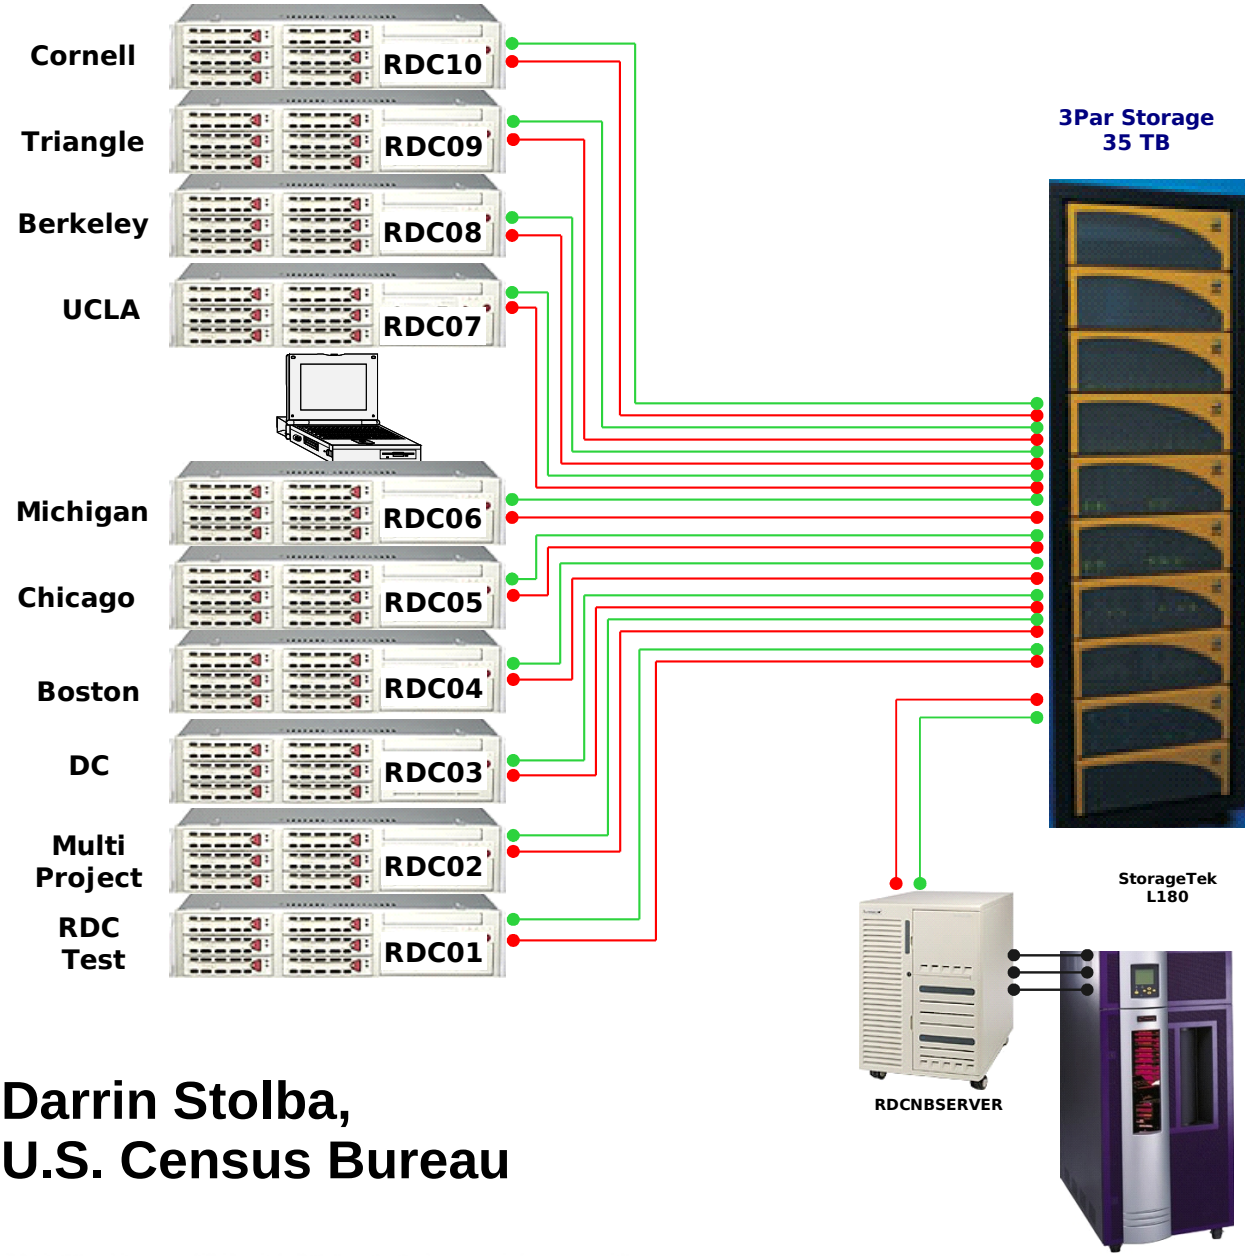

Darrin Stolba,
U.S. Census Bureau

RED HAT :: CHICAGO :: 2009 SUMMIT

Pros & Cons of Old Environment

- Pros
  - Local management of usage
    - The Gopher Effect (a.k.a. peer pressure)
  - Localized outages
- Cons
  - Administration of multiple servers
    - Updates (no internet access)
    - Security Checks
    - Where's Waldo?
  - Inefficient resource utilization (redundancy)
  - Localized outages

Specs of the New Blade System

- 6 IBM Blade Servers
  - 2-way dual-core
  - 16GB RAM
- Red Hat Enterprise Linux 5 (RHEL 5)
  - Linux Virtual Server (LVS)
  - Red Hat Global File System (GFS)
  - Red Hat Cluster Suite
  - Red Hat Virtualization Hypervisor
- All RDCs Log Into Same Cluster Alias (rdccluster2)

RED HAT :: CHICAGO :: 2009 SUMMIT

RDC
Virtual Machine Environment

3/25/2009

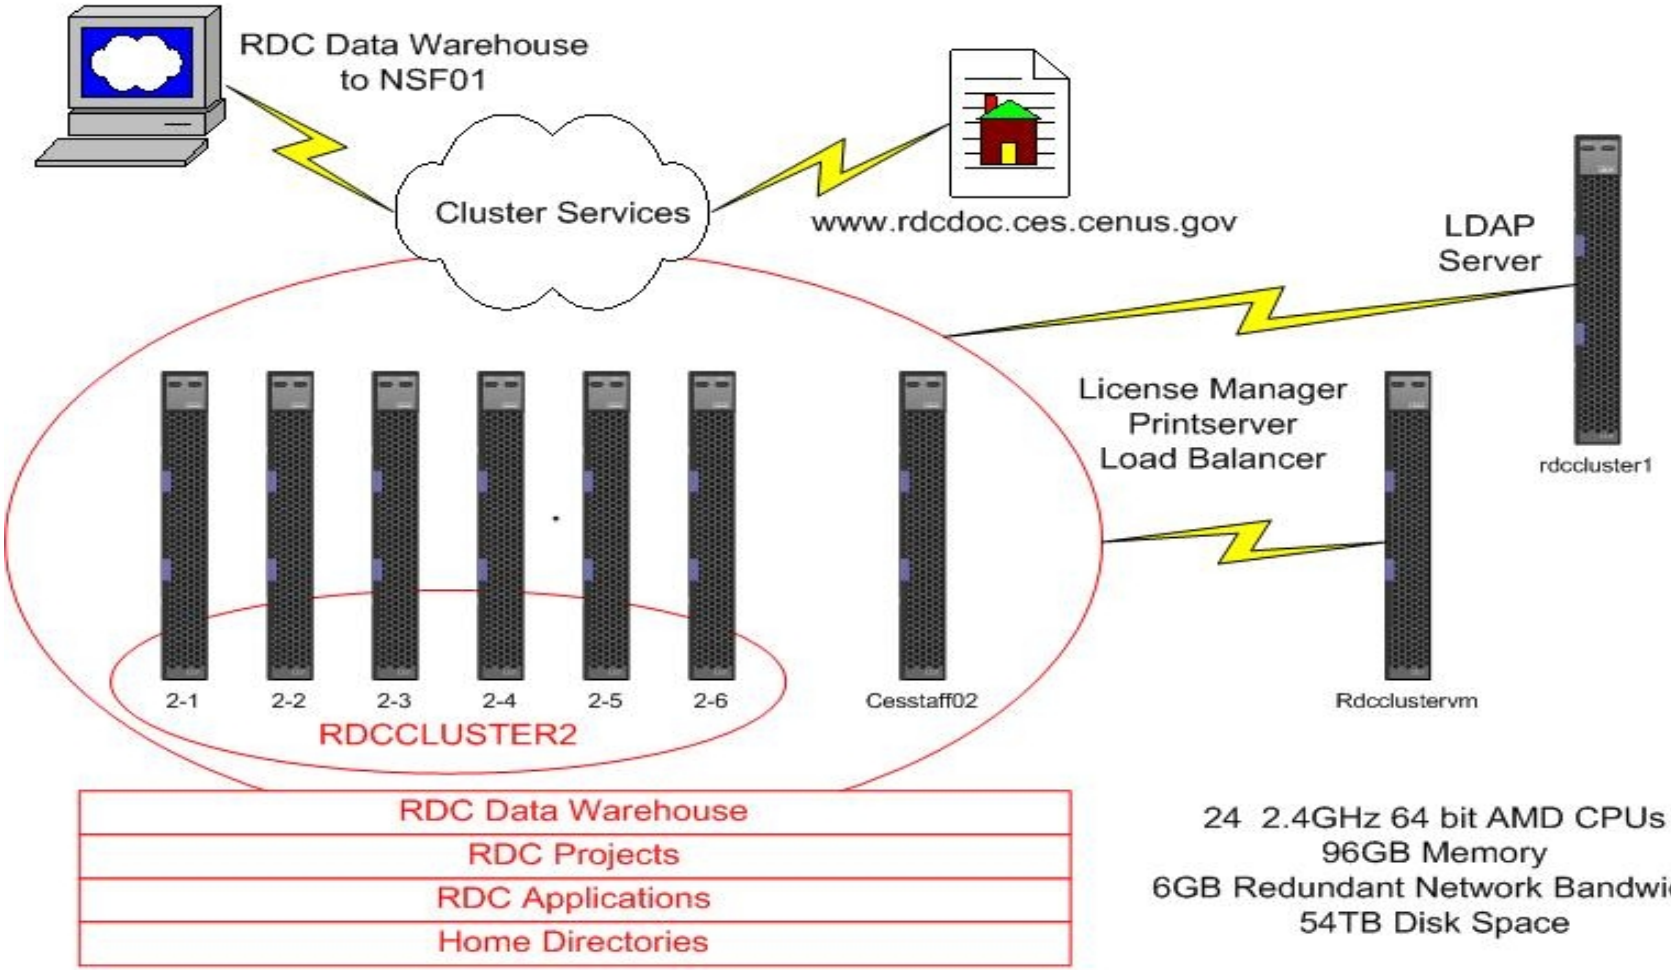

24 2.4GHz 64 bit AMD CPUs
96GB Memory
6GB Redundant Network Bandwidth
54TB Disk Space

Leo J Pleiman, Red Hat Professional Services

Benefits of Blade Servers

- Upgrade of existing system
  - Many servers hitting 100% (disk space and memory)
- More sharing capabilities
  - Shared space for staff
  - Spreads workload across all servers
  - 1 login for CES Staff
  - Building internal documentation Wiki
- Easier to maintain and expand
  - Shared software applications
  - LDAP maintains permissions/passwords
  - Expansion of RDC Network
- Better user experience

Potential Costs of Blades

- Resource hogs
  - Most users not used to shared environment
  - Lose local peer pressure effects
  - Train users on "good citizen" practices
- Network-wide outages

The Transition

- Testing, Testing, Testing
  - Critical to transfer user settings and permissions exactly (security)
  - Needed to ensure all files transferred from all servers (sporadic users)
  - Everyone has a favorite package
- Slow and steady conversion
  - Initial pilot by DC RDC
  - Move servers 1 by 1

Initial Success!

Success = Happy Users

- Jobs processed faster
- More disk space (especially temp space)
- Essentially seamless to users

Minor Challenges of Virtualization

- Managing user settings with LDAP
- Load Balancer puts users on “least busy” server
  - Makes tracking and troubleshooting difficult
- Residual processes and files clog servers
  - NX logins
  - Temp files
  - Hung/inactive processes

Build It And They Will Use It

- Balancing the load
  - Load balancing based on number of users on a node
  - 1 user can bring the node to a screeching halt
- Space Conservation
  - Hard to predict project space requirements
  - Difficult to get users to manage

Solutions So Far

- Development of log-in node
  - Use for non-CPU intense processes (e.g., editing)
  - Should solve some residue issues
  - Resource management software to manage the allocation of resources to other nodes
- Help the users
  - Improved documentation
  - *Good Citizen Practices* in each work space