

ARGONNE NATIONAL LABORATORY
9700 South Cass Avenue
Argonne, IL 60439

ANL/MCS-TM-267

The Production Cluster Construction Checklist

by

***Rémy Evard, Peter Beckman, Sandra Bittner, Richard Bradshaw, Susan Coghlan,
Narayan Desai, Brian Finley, Eugene Rackow, John-Paul Navarro***
{evard, beckman, bittner, bradshaw, smc, desai, finley, rackow, navarro}@mcs.anl.gov

Mathematics and Computer Science Division

Technical Memorandum No. 267

October 2003

This work was supported by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract W-31-109-ENG-38.

Argonne National Laboratory, with facilities in the states of Illinois and Idaho, is owned by the United States Government and operated by The University of Chicago under the provisions of a contract with the Department of Energy.

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor The University of Chicago, nor any of their employees or officers, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of document authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Available electronically at <http://www.doe.gov/bridge>

Available for a processing fee to U.S. Department of Energy and its contractors, in paper, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone: (865) 576-8401
fax: (865) 576-5728
email: reports@adonis.osti.gov

Contents

Introduction	1
1 Planning	2
1.1 Configuration Decisions	2
1.1.1 Requirements Gathering	2
1.1.2 Physical Constraints	2
1.1.3 Key Configuration Decisions	2
1.1.4 Management Approach	2
1.2 Vendor Input	2
2 Arrival	2
2.1 Installation Planning	2
2.2 Shipping	3
2.3 Inventory	3
3 Installation	3
3.1 Physical Installation	3
3.2 Initial Configuration	3
3.2.1 Base Networking Configurations	3
3.2.2 Hardware Control	3
3.2.3 Develop Minimal Base Image	3
3.2.4 Boot Process	3
3.2.5 Getting Images Out to All Computers	3
3.2.6 Remote Execution	4
3.3 System Check	4
3.3.1 Hardware Inventory and testing	4
3.3.2 Software Inventory	4
4 Configuration	4
4.1 Image Management	4
4.1.1 Grok the Configuration System	4
4.1.2 Define and Distribute Images	4
4.2 Basic Cluster Functions	4
4.2.1 Networks	4
4.2.2 Standard System Services	4
4.2.3 Configure File Systems	5
4.2.4 Cluster-specific Services	5
5 System Testing	5
5.1 Outage Tests	5
5.2 Component Tests	5
5.3 Benchmarking	5
5.4 Stress Tests	5
5.5 Application Tests	5
5.6 Acceptance Test	5

6	Preproduction	6
6.1	Timeline Check	6
6.2	Sanity Check.....	6
6.3	Final Configurations.....	6
6.4	Preparation for Users.....	6
6.4.1	Software for users.....	6
6.4.2	Scheduler and Access Policies.....	6
6.4.3	User Documentation.....	6
6.4.4	Account Management System	6
6.4.5	Allocation System.....	7
6.4.6	User Communication System	7
6.4.7	Trouble Reporting System.....	7
6.5	Early User Mode.....	7
7	Production	7
7.1	Announce Production	7
7.2	Ongoing Operation	7
7.3	Future Planning	7

The Production Cluster Construction Checklist

by

Rémy Evard, Peter Beckman, Sandra Bittner, Richard Bradshaw, Susan Coghlan,
Narayan Desai, Brian Finley, Eugene Rackow, John-Paul Navarro
{evard, beckman, bittner, bradshaw, smc, desai, finley, rackow, navarro}@mcs.anl.gov

Introduction

This document is a detailed checklist of the steps that one must go through to bring up a production computing cluster. The list starts with planning activities and culminates in the activities necessary to operate and sustain a production computing facility.

This checklist is derived from a number of experiences installing real-world, large-scale clusters. While each installation experience was unique, we were interested in determining the common characteristics across each deployment. We collected all of the to-do lists, presentations, notes, email messages, white board notes, and any other planning tools we could find from each of the installation activities. We combined them into a huge, messy diagram that was probably impossible to understand without having been involved in its creation but was excellent for identifying differences and commonalities. After organizing, checking, and distilling the information, we created the checklist presented here.

Interesting is the fact that the high-level activities on the resulting list are neither cluster nor computer specific. Most of these activities would be followed when installing a production computer of any architecture or when installing any kind of complex facility that will eventually support users.

The purpose of this list is not to give step-by-step instructions but rather to serve as a guide and a reminder. The items on the list are necessarily brief statements. Detailed explanations of these would go beyond the intended scope of the list.

The list is organized in outline fashion. The major phases of construction are individual sections. Each of the subsections is a task or subtask in that phase.

The items on this list are presented in a logical sequence, in approximately the order that one would follow if one were to start with a budget and an idea. However, every cluster is different, and every situation for using clusters is different. Most likely, no one would ever follow the steps here in this exact order; many things can be done in a different order, simultaneously, or skipped altogether. The list, for example, may place more emphasis on testing than many sites formally will.

Please send recommendations for improvements to future versions of this list to Rémy Evard at evard@mcs.anl.gov.

1 Planning

Having determined that you need a cluster and have a budget, determine specifically what it is that you want and who will provide it. If working with a vendor, your goal is to have a final quote and contract.

1.1 Configuration Decisions

Develop a detailed plan for the cluster.

1.1.1 Requirements Gathering

- User expectations
- Usage model
- Success criteria

1.1.2 Physical Constraints

- Space limitations
- Power limitations
- Cooling limitations

1.1.3 Key Configuration Decisions

- Node architecture
- Network architecture
- Storage and I/O architecture
- Special devices (e.g. visualization, acquisition)

1.1.4 Management Approach

- Core administrative functionality
- Security policies
- Software model

1.2 Vendor Input

Work with vendors to define a system that meets your requirements and fits in your budget. Most likely this will be an iterative process in which you adapt your configuration based on what can be provided.

- To vendor or not to vendor – determine whether you will be working with a systems integrator or rolling your own cluster from parts.
- RFP
- Quotes
- Negotiation
- Contract

2 Arrival

Move from having a final quote from the vendor to having all of the correct hardware at your site, ready to be installed.

2.1 Installation Planning

- Confirmation of vendor's role in installation and configuration
- Floor plans
- Rack layout plans

- Wiring plans
- Layout logistics
- Power availability and connectivity
- Fire suppression

2.2 Shipping

- Confirmation of date and shipping logistics
- Unloading and unpacking space

2.3 Inventory

- Check shipping receipts against boxes that have arrived
- Visually inspect boxes for damage
- Check contents of boxes against final quote from vendor

3 Installation

Set up the system, and carry out initial inventory tests.

3.1 Physical Installation

Install the hardware in the racks, attach the networks and storage devices, and connect everything to the power system. If a system integrator is involved, it will do all of this but may need supervision.

3.2 Initial Configuration

Get the cluster to the point where it is on and responsive enough that you can run simple jobs across the entire cluster in order to do system checking.

3.2.1 Base Networking Configurations

- Have a network space and network name plan
- Prepare the network switch(es)
- Implement the network mechanisms in the base image or boot mechanisms

3.2.2 Hardware Control

- Set up the remote power control mechanism, if any
- Set up the remote console mechanism, if any
- Force hardware reboots on individual components to test them

3.2.3 Develop Minimal Base Image

This may be provided by the vendor, may be minimal, or may be nearly complete.

3.2.4 Boot Process

- Set up the boot mechanism
- If using DHCP, may need to collect Ethernet addresses
- Set up the boot server with appropriate images

3.2.5 Getting Images Out to All Computers

- Boot all of the nodes and servers, ensuring that they get the correct images

3.2.6 Remote Execution

- Determine how to execute commands on all systems

3.3 System Check

With the system connected and responsive enough to boot all nodes and execute programs, it is now possible to carry out initial functionality and inventory tests.

3.3.1 Hardware Inventory and testing

Confirm that each of the components has the right parts and that the system is connected properly.

- Check CPUs, memory, disk, and temperature of systems.
- Check pings across the networks.
- Test storage accessibility, configuration, and size.

3.3.2 Software Inventory

When using software provided by a third party such as with prebuilt vendor systems or HPC distributions, check to see what software exists.

- Base OS, libraries, kernels, etc.
- What else is installed? Does it look sane?

4 Configuration

Carry out all of the systems administration tasks that transform the individual hardware components into a functioning cluster.

4.1 Image Management

4.1.1 Grok the Configuration System

Understand the image and configuration mechanisms. Either the system provided will have certain expectations and tools, or you will need to determine how to do this yourself.

- How do you define images?
- How do you push out images?
- How do you push out changes?
- How do you install software?
- How do you install patches?

4.1.2 Define and Distribute Images

Build the base images for all parts of the cluster, modified as you see fit.

- Primary cluster server
- Any secondary servers: login, management, other
- Compute nodes
- Any specialized nodes

4.2 Basic Cluster Functions

4.2.1 Networks

- Tune or modify the IP space and Ethernet
- Configure any other networks

4.2.2 Standard System Services

- Configure all the services you'll be using: DNS, NTP, SSH/RSH, NIS, any others

- Disable all the services you won't be using

4.2.3 Configure File Systems

- File systems on individual nodes
- Home file systems and servers
- Cluster-wide file systems and servers
- Parallel file systems and servers
- Other specialized I/O such as mass storage access

4.2.4 Cluster-specific Services

- Schedulers
- Resource managers
- Parallel tools

5 System Testing

With the cluster basically operational, carry out tests to confirm that the system is working correctly. Ensure that you have a strategy for addressing any problems that you discover.

5.1 Outage Tests

- The reboot test. Reboot the entire system, see if it comes back completely without help.
- The power lossage test. If you have the nerve, test to see what happens if you kill the power.

5.2 Component Tests

Test all individual system components in the cluster.

- File system performance tests.
- The “everybody core dump at once” test
- High availability file system tests. What happens if a server goes down?
- Network connectivity tests
- Network performance tests
- File transfer into and out of the cluster

5.3 Benchmarking

Run standard benchmarks such as LINPACK and the NAS benchmarks to test overall system performance.

5.4 Stress Tests

Run any stress tests of individual components or the overall system that you deem necessary. Typically these are tough on the system, run for a long time, or both.

5.5 Application Tests

If certain applications are essential to your user community, run these applications to test functionality, performance, and correctness.

5.6 Acceptance Test

Run acceptance tests, typically as determined during the contract process, to determine what the vendor must address before the system is accepted.

6 Preproduction

At this point, the system is known to be working. Complete the work necessary to support users.

6.1 Timeline Check

Make or confirm the plan to shift into production mode. The timeline will drive the priorities for software installation and user support.

6.2 Sanity Check

- Confirm ability to do core systems administration tasks:
 - Add and remove nodes
 - Replace hardware
 - Add, upgrade, and remove software
 - Modify and distribute images
 - Remotely manage hardware
 - Add and remove user accounts
- Confirm usage policies
- Confirm security implementation

6.3 Final Configurations

Configure and test any remaining systems.

- Monitoring systems
- Backups
- Failover services for DNS, NIS, etc
- Check all the early to-do lists for lingering issues

6.4 Preparation for Users

Install the software and systems necessary to support and manage users on the system.

6.4.1 Software for users

- Compilers, development libraries, debuggers.
- Third-party software

6.4.2 Scheduler and Access Policies

- Develop initial scheduler policies
- Confirm node access policies, i.e., that computing nodes are only available to scheduled users.
- Allocate certain systems for special use such as interactive development and testing.

6.4.3 User Documentation

- FAQs such as “how to get an account” and “how to report problems”
- Documentation specific to the system, such as hostnames, access policies, etc.
- Quick-start guides
- Tutorials
- Software documentation

6.4.4 Account Management System

- A mechanism for users to request accounts
- A mechanism for creating users accounts
- Paper trails as necessary for your organization

6.4.5 Allocation System

- If usage will be tracked, determine how users will request time on the system, how allocations will be enforced, and how usage will be monitored

6.4.6 User Communication System

- Set up systems for sending announcements to users
- Consider mailing lists for user discussion

6.4.7 Trouble Reporting System

- Set up systems for users to report problems, typically a trouble ticket system

6.5 Early User Mode

- Determine criteria for early users
- Set expectations for system availability and reliability
- Consider time (duration, schedule) for regular system downtime
- Respond to user issues
- Meet with early users

7 Production

Having completed all the advance preparation, open the system up to all appropriate users. Shift emphasis from configuration to operation. Most of the items described here will be carried out simultaneously.

7.1 Announce Production

- Choose a date to change to production.
- Hold any necessary ceremonies, announcements, etc.

7.2 Ongoing Operation

Operations mode is not oriented around a specific list as much as continued support of the user community, the system, and striving to achieve the goal of the system. With a large user community, the workload will be about the same as it was for system configuration.

- Support users
- Respond to problems
- Respond to emergencies
- Tune and upgrade the system
- Continue regular tasks
- Create documentation
- Other – very much subject to site expectations

7.3 Future Planning

Remember to learn from experiences and plan for the future changes to the system or for future systems.