

Microway[®] Cluster Manual

including Microway MPI Link-Checker[™] and
MCMS[™] with IPMImon[™]

Copyright 2023
by
Microway, Inc.



Microway, the Microway logo, NumberSmasher, WhisperStation, Navion, IPMImon, MCMS, Microway MPI Diagnostic Suite, MPI Link-Checker are trademarks or registered trademarks of Microway, Inc. AMD, AMD Athlon, Opteron, EPYC, Ryzen, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. in the United States and other countries. Intel, the Intel logo, Xeon, and Phi are trademarks or registered trademarks of Intel Corporation, or its subsidiaries, in the United States and other countries. Mellanox is a trademark of Mellanox Technologies, Inc. in the United States and other countries. All other names and trademarks are owned by their respective companies.

Table of Contents

Introduction.....	5
Microway Technical Support Contacts.....	5
Microway Enclosed Cluster Letter.....	5
Quick Start Guide.....	5
Assemble Hardware.....	5
Connect Network.....	5
Power on Dedicated Storage systems.....	6
Power on Head Node.....	6
Power on Compute Nodes.....	6
Microway Cluster Management Software (MCMS™).....	6
Message Passing Interface (MPI).....	7
Lmod.....	7
Slurm.....	8
Administration Scripts.....	8
Important Cluster Configuration.....	9
Change Passwords.....	9
IPMImon™ Alerts.....	9
Other Hardware Alerts.....	9
External Network.....	9
Cluster Hardware.....	12
Environmental Requirements.....	12
Unpacking and Assembly Instructions.....	12
Cable Routing.....	13
Power On Cluster.....	14
Storage Administration.....	15
Individual Hard Drives.....	15
Software RAID.....	15
ZFS.....	15
RAID/NAS Systems.....	15
IPMImon™ Hardware Monitoring and Control.....	16
Cluster Software.....	17
Linux Configuration.....	17
Introduction.....	17
Environment Variables.....	17
MPI-Selector.....	17
Lmod.....	17
System Services.....	18
Configuration Files.....	19
Important Linux Kernel Upgrade Note.....	20
Microway Cluster Scripts.....	21
Microway Cluster Management Software (MCMS™).....	23
Introduction.....	23
Integration with IPMImon™.....	23
Node Backup and Restore.....	23
LDAP Configuration.....	23
LDAP CLI management.....	24
Getting Started:.....	25
Cluster Control.....	26
Account Administration.....	29
Node Backup.....	30
Node Restore.....	33
Other User Administration Issues.....	34

MCMS Administrative User Configuration.....	34
Grafana.....	35
MCMS Installation.....	38
MCMS-network Installation and configuration.....	38
MCMS head and client package installation.....	42
IPMImon™ Software Configuration.....	43
IPMImon Software Component Overview.....	43
Automatic Email Notification, Shutdown and Power Off.....	43
IPMImon Daemon Control.....	45
Microway HPC software.....	46
Compilers.....	46
Gnu Compiler Collection (GCC).....	46
Intel® oneAPI Compiler Suite.....	46
NVIDIA HPC SDK (formerly known as PGI® Compiler).....	46
Message Passing Interface (MPI).....	47
C Example.....	47
Fortran Example.....	47
MPI Implementations.....	47
OpenMPI.....	47
MVAPICH2.....	47
MPICH4.....	48
Intel MPI.....	48
InfiniBand and OmniPath Interconnects.....	48
Important InfiniBand Performance Note.....	48
Slurm.....	48
Microway MPI Diagnostic Suite (MMDS).....	51
MPI Link-Checker™.....	51
Microway Infiniscope™.....	52
Cluster Extensions.....	54
Performance Tips and Cluster Tuning.....	55
Compiler Optimization.....	55
Network Tuning.....	55
Extraneous Services.....	55
Troubleshooting.....	56
Linux Stress Tests.....	56
DOA (Dead On Arrival).....	56
System Does Not Power On.....	56
System Powers On, but Does Not Boot.....	56
Network Connectivity.....	56
System Crash.....	57
Hard Drive Errors.....	57
RAID Failure.....	57
Bad Memory.....	59
EDAC reporting.....	59
MCELOG reporting.....	59
Corrected vs Uncorrected errors.....	60
Isolating the problem by swapping components.....	60
Bad CPU.....	60
MPI Problems.....	61
Problems Starting MPI Jobs.....	61
MPI Jobs Exit Prematurely.....	61
Basic troubleshooting rules to live by.....	62

Introduction

Thank you for purchasing a Microway cluster! This manual will help ensure that you get the most from your investment. It contains information and instructions for installing the systems, using your cluster and troubleshooting any problems. We highly recommend that you take the time to read this information.

You may find that certain portions of this manual do not exactly match the cluster that was shipped to you. Because Microway prides itself on configuring clusters to customer's exact specifications, we may have changed items to ensure the cluster would operate as you requested. If you have any questions or concerns that are not addressed here, please contact Microway's expert technical support.

We hope this cluster will meet all of your computing needs!

Microway Technical Support Contacts

Please contact our technical support staff if you should have any questions or problems. You can reach us by phone during regular business hours Eastern Standard Time. Please have your user (invoice) number or system serial number handy so that we may better serve you.

Microway, Inc.
12 Richards Road
Plymouth, MA 02360

Telephone: (508) 746-7341
Fax: (508) 746-4678
tech@microway.com
<http://www.microway.com>

Microway Enclosed Cluster Letter

A single 8.5" x 11" white envelope was included with your cluster - usually attached to the head node. Inside, you will find a short letter which contains information that is unique to your cluster. You will need to consult this documentation while reading the manual to ensure you properly configure your cluster. It includes passwords, directions for the proper placement of systems in the cabinet, and the correct cabling of the head node and compute nodes.

Quick Start Guide

If you would like to immediately begin using your cluster, please read this brief overview. The *Quick Start Guide* is not intended to be a step-by-step walk-through for cluster installation and configuration. For more details on setting up your cluster, please skim this guide and then refer to the *Cluster Hardware* section.

Assemble Hardware

Unpack the cabinet and move it into position. Unpack the cluster components and install them into the cabinet. You will need to refer to the letter that Microway enclosed to determine where cluster components should be installed in the cabinet. Connect power cables to all components, taking note of the power requirements of each and appropriately load balancing them on the power strips. More detailed assembly instructions can be found in the *Cluster Hardware* section below.

Connect Network

To satisfy the cluster software, you will need to connect each compute node and the head node to the provided ethernet switch. Please note that both the head node and the compute nodes may contain more than one ethernet port. You will need to refer to the enclosed letter from Microway to determine which port(s) should connect to the network switch.

Power on Dedicated Storage systems

If you have dedicated storage servers in your cluster, please power those on first. If those systems utilize external JBOD chassis, ensure those chassis are powered on and connected prior to powering the system on.

Power on Head Node

If you have any external storage JBODs, power those on first. The first system to be powered on after any external JBODs is the head node. A series of computer tests will appear on the screen, followed by boot messages, as the system loads the Linux operating system. Note that these systems have more CPUs and ECC memory than a typical desktop computer and may take up to 90 seconds before video is displayed.

Once booted, the system will ask you for a user login and password. A test user “**microway**” has been configured with a password of “**microway**”. Your root password has also been set to “**microway**”. Clusters running Ubuntu typically do not have a root password set and instead rely on sudo usage; however, if they have LSI RAID management software installed the root password will be set as those tools require root login for full functionality.

Power on Compute Nodes

Be sure that you have logged on to the head node before powering on your compute nodes to ensure that all services have started on the head node. To power on your compute nodes, simply press each on/off switch on the front of the chassis. You can also power all nodes on from within the MCMS web gui. The systems are preset to automatically boot into Linux after a short system/memory check. A quick network check to assure connection to all nodes can be done by running `do.ssh` as root on the head node. If you have any trouble, refer to the *Troubleshooting* section of this manual before continuing.

Microway Cluster Management Software (MCMS™)

MCMS may be installed on your cluster. To view the status of your cluster from the head node, double click the MCMS icon on the desktop. You can also open a web browser and use the url <https://localhost/mcms/>. To view the status from another computer on your company or campus network, use the head node's external domain name i.e. https://external_hostname.yourcompany.com/mcms/

Microway has configured apache with a self-signed certificate for SSL. The certificate can be replaced with an official CA cert following standard procedures for your distribution if desired. Please see the section on mcms-certgen for instructions on managing the cluster's SSL certificates.

The MCMS interface starts with a login page. User accounts designated as MCMS administrators may login with their regular username and password. A fresh installation provides the user “microway” as an admin by default. See the MCMS user administration section for details on management.

Current page: Login page

Cluster Accounts Help Logout

User: Password:

Log In

MCMS™ - Microway Cluster Management Software

Once logged in, you will be presented with a graphical interface to copy files, run commands, add/remove/modify users, perform node backups and restores, and control reboots and shutdowns on either a particular node or the complete cluster. If your cluster includes IPMI controllers, Microway's NodeWatch feature provides a table with all sensor readings. The groupings of Fans, Volts, and Temps can be collapsed into a status light by clicking on the column heading. The color coding uses green for healthy, amber for alert, and red for critical. The image below shows healthy values for all metrics and has the Fans column expanded to show full details.

Current page: Cluster Control

Cluster Accounts Help Logout

Copy file from head to node(s)

From: To path: All

Perform a command on node(s) as root

Command: All

Shutdown node(s):

Reboot node(s):

NodeWatch™

Nodes		Fans				Volts	Temps	OS
Name	FAN3	FAN5	FAN4	FAN6	Status	Status	Gstat	
head	1900	1000	1000	1700				
node1	1800	900	800	1500				
node2	1800	900	900	1600				
node3	1700	800	800	1400				
node4	1700	800	800	1500				

MCMS™ - Microway Cluster Management Software

Please see the MCMS manual section for full details on the use and configuration of the suite.

Message Passing Interface (MPI)

If you requested an MPI library, the base installation directory for all MPI files is `/usr/local/mpi` and is owned by root. To run an MPI example, login as microway and from the home directory run:

```
./nas-test.sh
```

This simple MPI test script will run an appropriately sized set of tests from the NAS Parallel Benchmarks suite in a loop. You can hit Ctrl-C to quit the test. If you requested a cluster with alternate HPC libraries, MPI may not be installed on your cluster.

When installed with multiple MPI library versions, Microway clusters will be deployed with either `mpi-selector` or `Lmod` (commonly referred to as “environment modules”) to enable choosing the library version you wish to use. See the cluster software section of the manual for details on their use. The “`module`” command is used to control which `Lmod` environment modules are loaded. See the `Lmod` section for details on usage if your cluster uses `Lmod` instead of `mpi-selector`.

Lmod

If included in your cluster, `Lmod` may be used to manage multiple versions of libraries, compilers, and other similar software where multiple options may be available. The command “`module`” can be used to manage your loaded environment modules. See the `Lmod` section in cluster software for more information.

Slurm

The SLURM scheduler is included in most Microway cluster deployments. It provides job scheduling and queuing to enable you to effectively share your cluster resources.

Depending on the distribution used and the Slurm features required for your cluster, Slurm will either be installed from a package repository or built into packages from source. You can check the details of the installed packages with your package manager to determine the method used.

The head node runs the `slurmctld` service. All nodes including the head run the `slurmd` service. These services can be controlled with `systemctl` commands. Example batch submission scripts are in the home directory for the test user `microway`. You can setup node and partition information in the file `/etc/slurm/slurm.conf`. This file must be the same on all systems. If you modify it on the head node, run `"scp /etc/slurm/slurm.conf /etc/slurm.conf"` as root to distribute the updated file to all systems. Then run `"systemctl restart slurmctld"` to restart the controller daemon on the head node and finally `"scom -parallel systemctl restart slurmd"` to restart the main Slurm daemon on all systems. You can find out more at <https://slurm.schedmd.com/>

Administration Scripts

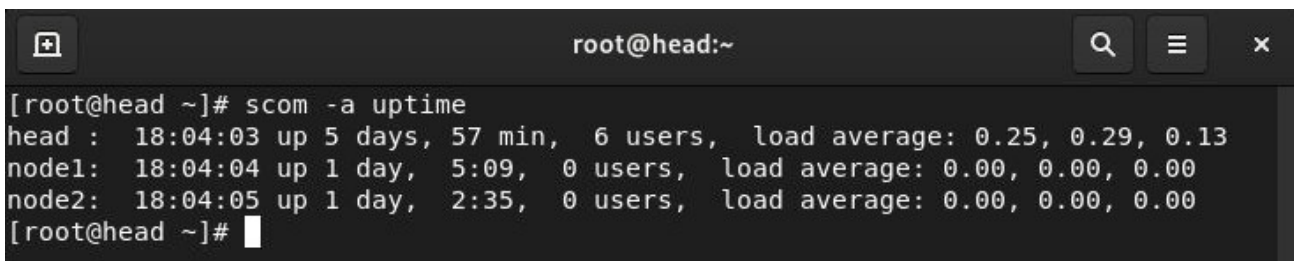
Additional utilities that one may find useful are scripts `scp` and `scom`. The scripts use `scp` and `ssh` to copy files throughout the cluster and to run commands throughout the cluster.

The `scp` utility will copy a file to the head node and all of the compute nodes on the cluster. For example, to copy an updated `/etc/hosts` file to the head node and all compute nodes run:

```
scp /path/to/new/etc/hosts /etc/hosts
```

The `scom` command will execute a command on the head node and all of the compute nodes. For example, to check uptime on all nodes

```
scom -a uptime
```



```
root@head:~  
[root@head ~]# scom -a uptime  
head : 18:04:03 up 5 days, 57 min, 6 users, load average: 0.25, 0.29, 0.13  
node1: 18:04:04 up 1 day, 5:09, 0 users, load average: 0.00, 0.00, 0.00  
node2: 18:04:05 up 1 day, 2:35, 0 users, load average: 0.00, 0.00, 0.00  
[root@head ~]#
```

The `-a` switch tells `scom` to print the node name on each line of output.

Please keep in mind that `scp` and `scom` operate on the head node and the compute nodes. Take care not to overwrite the head node's configuration with configurations from the compute nodes, as the head node's configuration is unique for some services. There are additional variants of the scripts with `-nodes` or `-parallel` as part of the name. `-nodes` variants work only on the compute nodes (exclude head node) and `-parallel` versions work on all systems simultaneously but do not guarantee ordered output. The parallel versions provide a `"-w"` parameter as well that wait for all commands to complete before returning to the command prompt.

For more details and other useful scripts, please refer to the *Cluster Software* section of this manual.

To power off your cluster, first shutdown all nodes by running `shutdown_cluster` from the head node. Then shutdown the head node by running `poweroff`. Similarly, running `reboot_cluster` will reboot all nodes.

Important Cluster Configuration

Before beginning to use your cluster, there are several items that need to be addressed. These items will improve the security of the cluster and simplify administration down the road.

Change Passwords

The default passwords that shipped with the cluster should be immediately changed to prevent unauthorized tampering. Open the *MCMS Account Administration* panel and change the password for each user. If MCMS is not installed, login as root and run “passwd”. Then change the microway password by running “passwd microway”. If you change your root password through MCMS it will be changed on all systems. If you change it from the command line manually you must change it on all compute nodes as well.

If the system has a RAID controller, you will need to change the RAID administrators password. This user has permissions to create, modify and even destroy the storage array.

For Areca controllers, open a web browser and go to <http://localhost:81>. The default login is admin with password 0000. On Linux systems you can also do this at the command line with the utility cli64. Run “help” within the utility for information on usage.

For LSI (Broadcom/Avago) controllers use the webgui LSA, or “LSI Storage Authority”. LSA can be connected to either via the desktop icon provided, or by opening “<http://localhost:2463>” in your web browser. Log in with the root user and password. Alternatively you can manage your LSI controller through the command line using the storcli64 command as root. We encourage you to read the man pages and use your favorite search engine on the Internet to learn more about these commands.

Managed network switches also have usernames and passwords to secure. Unless otherwise noted in your custom letter accompanying your order, we have not set any passwords for these devices. Please refer to the manufacturer's documentation for specific instructions on changing passwords. The switch can usually be managed using the web interface provided by the switch. The dhcp log or lease file on your head node can indicate which IP was assigned to the switch. You may use your web browser to pull up this IP and manage the switch. You may also set a static IP in the web interface and add it to your /etc/hosts file to enable connecting to it by name. In certain cases Microway will preset a static IP on a cluster switch and add it to /etc/hosts for configuration purposes. Many switch management interfaces do not save changes permanently by default. **Be sure to store your new settings in the switch's permanent memory!** If you do not, your settings will be erased if the switch loses power or is reset.

IPMI™ Alerts

If your cluster is equipped with IPMI controllers for our remote hardware monitoring solution, you can receive email alerts whenever fan, voltage, or temperature measurements exceed safe values. Additionally, they can automatically shutdown and power off any nodes which exceed these values. **By default, automatic email notification and power off are disabled!** Please see the *IPMI™ Software Configuration* section for installation and configuration details.

Other Hardware Alerts

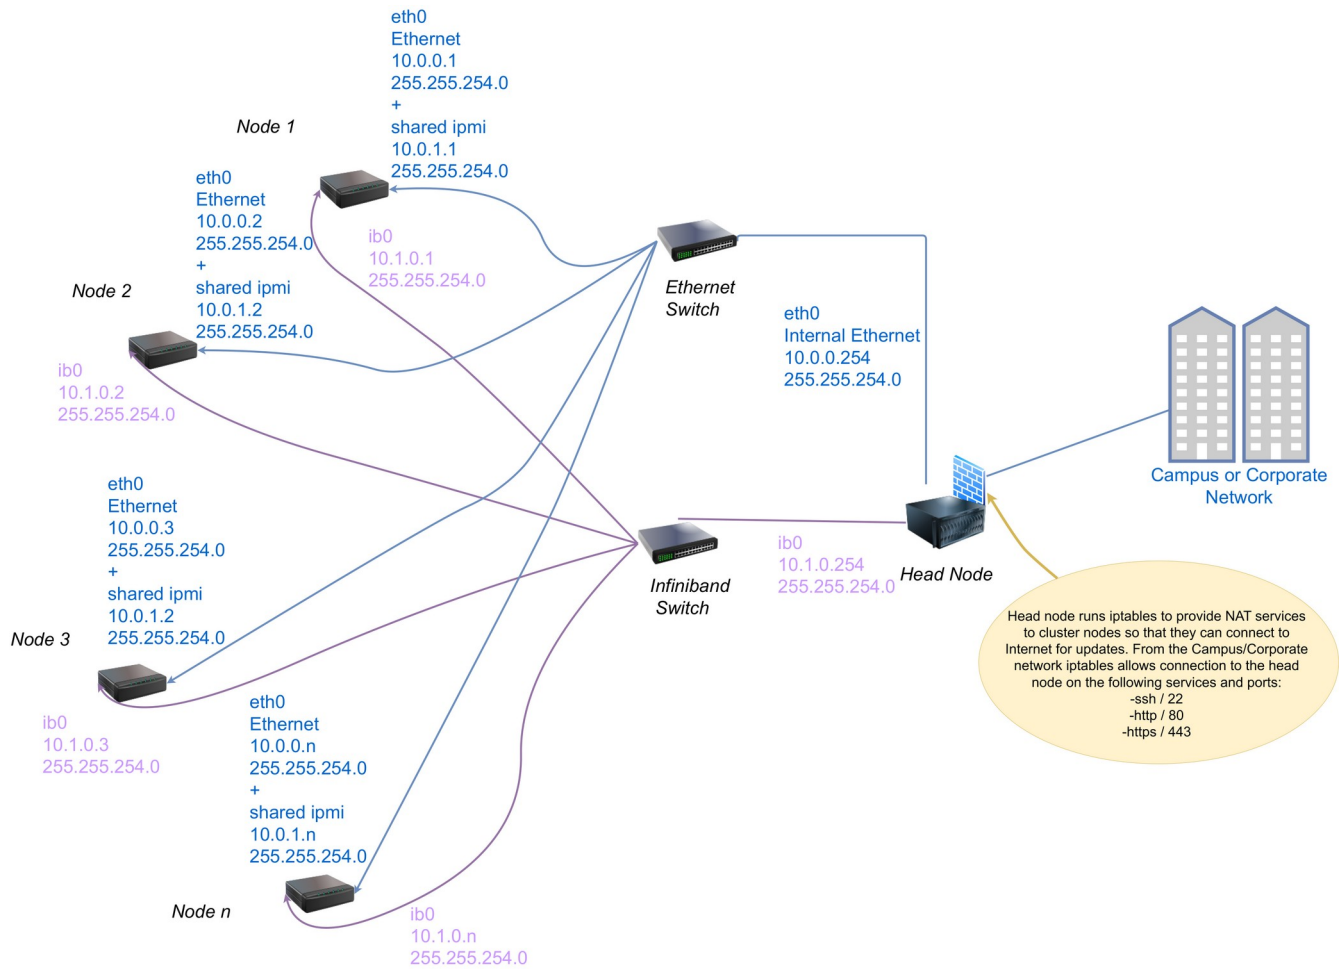
While changing the passwords on your RAID controller(s), investigate the email alert capabilities. Most RAID controllers will send alerts to administrators in case of array degradation or failure. Be sure to enable this feature, as hard drive failures will go unnoticed and lead to data loss. See the *Storage Administration* section for additional information and the *Troubleshooting* section for details on finding disk failures and replacing disks.

External Network

If you specified your head node's external network configuration on the Microway Software Questionnaire, it will have already been set for you. If not, you will need to set it now to enable communication with the rest of your organization's network. **However, please do not change the host names or domain names of your cluster!** See the *Linux Configuration* section for more details and instructions for changing the network configuration.

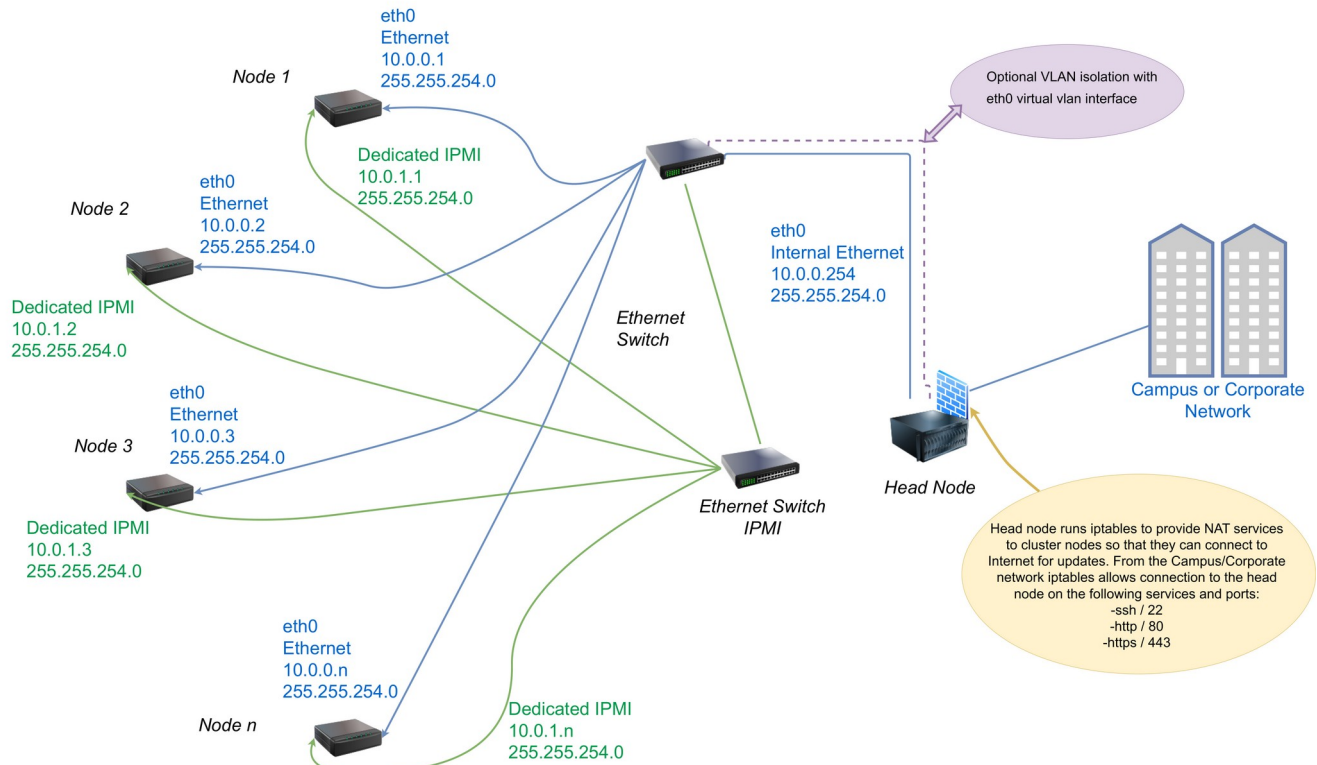
Network Diagrams

Cluster with Infiniband



Example network diagram using shared IPMI and InfiniBand

Cluster with Dedicated IPMI



Example network diagram including dedicated IPMI connections. The diagram illustrates an optional VLAN interface on the head node that can be used to isolate IPMI access.

Cluster Hardware

Environmental Requirements

Microway clusters require the same cool, dry environment as other computer equipment. Your cluster will produce a significant amount of heat, so you should consult your facility's Maintenance/HVAC department for requirements specific to your unique cluster.

It is your responsibility to verify that the air entering the cabinet at all elevations is cold enough to cool the cluster. If you own an IPMI enabled system, this can easily be verified by observing the temps for the various systems in the cabinet. Periodically recording the temperatures and fan speeds of the nodes will make it possible to detect air flow blockages that may develop within the nodes over time due to dust.

It only takes a few minutes to damage the nodes of a cluster if the cluster is allowed to run in a room that is inadequately cooled. If your cluster has backup power, but your AC does not, you must initiate a shutdown the moment your data center loses power. It is also important to monitor the status of your AC and to initiate a shutdown anytime the room or data center temperatures exceed safe levels. Waiting for human intervention in either cooling failure scenario can quickly result in cluster damage. Automatic shutdowns should be enabled with IPMImon to help prevent this possibility.

A second method for detecting an AC failure in a room is required, and must be installed in such a manner that if the AC fails, a cluster shutdown is initiated. It is the user's responsibility to periodically test this AC failure shutdown mechanism to verify that it works correctly.

Also note that high speed fans are required to cool these dense HPC systems. The cluster will create a fair amount of noise, and should not be positioned near anyone's work area. Sitting next to a cabinet of nodes for an extended period of time may cause hearing damage. Persons working near the cluster should wear headsets or earplugs for protection.

Unpacking and Assembly Instructions

Remove the large cabinet from its pallet. Set the cabinet in the upright position, move it to its final location and lower the stabilizer bars at the bottom to immobilize the unit. Remove both the front and rear door to the cabinet. The rails for some chassis will have tie wraps fastening them into place during shipment. Cabinet keys will be similarly tie-wrapped onto the rails. Remove these tie wraps from the front of each slide rail. Some systems do not contain inner slide rails. These units simply slide into the existing outer rails mounted in the cabinet.

Depending on the complexity of your cluster, the documentation included with your cluster letter may include a diagram for the layout of your cluster in the cabinet. You will find that Microway has already installed the rails in your cabinet to match this layout. Follow this guide for an efficient installation process.

UPS units should always be located in the bottom-most position of the cabinet to accommodate their weight and for ease of access to outside power connections. Be sure to install the UPS before any other components as it is very difficult to position the UPS once the cabinet is full.

Locate and unpack the system to be installed at the bottom of the cabinet. Slide out the inner portion of the rail, if any, until it snaps into place and then insert the system into the rail. Continue this procedure for the next consecutive node number. Each system has a node number located on the outside packing box as well as on the rear or side of the actual computer chassis. Going in consecutive order during installation will keep the cluster well organized should one need to pull a node for any reason. Once the systems are installed in the cabinet, secure them by either tightening the thumb screws on the front of the chassis or by screwing down chassis ears, depending on the style of the chassis.

Install all included RAID arrays, tape drives, network switches, KVM switches and LCD monitor. Refer to the cabinet layout included in the enclosed letter from Microway to determine their positions.

Cable Routing

Looking at the front of the 19 inch rack cabinet, there are panels on each side of the rack cabinet that, if removed, allow easy routing and tie-up of the necessary network cables connected to the rear of the systems.

Route and secure cables as shown in *Figures 1 & 2*. Refer to the enclosed letter from Microway to determine which network ports on the head node and compute nodes should be used. Some clusters may use dedicated IPMI connections. The letter accompanying your order will indicate these details.

Bundling the cables in groups of 3 or 4 and attaching to the inside frame of the cabinet works well. The routing and securing of network cables in this document are only suggestions from Microway to aid in network cable



Figure 1

management. Securing the cables as documented will allow more airflow and make it easier to remove nodes.



Figure 2

For proper power distribution, all power strip plugs need to be routed out of the of the cabinet and connected to a known good power source. The power consumption of HPC systems can be quite high in many cases. Depending on the CPUs, GPUs, storage, and other components, a system could consume anywhere between 300-5000 Watts. Power must be distributed across PDU banks to ensure proper operation. Available power per bank is calculated by multiplying the voltage by the amperage and then reducing by 20% (multiply by 0.8). A 20A/208V bank has 3328W of usable sustained power. A 15A/115V bank only provides 1440W usable. Some systems require 200-240V to operate and will not function with 110-120V power. Refer to your Microway order for details including the estimated maximum power consumption per system. The labels on your power supply provide information about the supported input voltages and the maximum power the supply can handle but does not necessary reflect the actual consumption of the system as built. Actual power consumption is typically considerably lower than the maximum the power supply is capable of.

Depending on the size of the cluster and the power configuration required, a variety of power strips may be included. In most cabinets there will be a single power strip on each side of the cabinet providing several banks of outlets each. You must make sure to balance power such that you are not drawing too much from any one bank. In the case of UPS connections, the systems connect to the cabinet power strips and the cabinet power strips connect to one of the UPS plugs. Some configurations will plug directly into UPS units, but that is not typical. Please note that many UPS units contain several separate power banks. The power ratings of each bank is printed next to the power plug. Be sure to adhere to the ratings for each bank when making the power connections. Some models can support 15A per bank while others support more. The industry standard is to connect a maximum load of 80% of a breaker's rating. Most PDUs and UPS devices will provide this value in their literature. For a 20A circuit this would allow connecting equipment drawing a constant 16A.

If your cluster includes an InfiniBand network, you will need to connect those cables to the InfiniBand switch(es). Note that InfiniBand cables are bulkier than Ethernet cables, so you will need to take care when bundling InfiniBand cables. If you have multiple switches, please refer to the letter accompanying your cluster for details on connecting them correctly. InfiniBand cables should make a gentle click when fully inserted. If you do not feel this, the cable may be upside down. The cables have a release tab that must be pulled to remove the cable. If the cable is removable without using the release tab, it is likely upside down.

If a KVM was ordered with your cluster, connect the desired systems to the KVM switch, running the cabling along the sides of the cabinet. You may find it easier to bundle the network cabling on one side of the cabinet and the KVM cabling on the other.

After cables are routed and secured, replace the cabinet side panel(s) and doors.

Power On Cluster

Because different portions of your cluster rely on each other, you should not power on your cluster by turning everything on at once. Power on items in this order (when powering off the cluster, reverse this list):

1. Network, InfiniBand and KVM Switches. Some switches (ie Cisco) take a very long time to initialize so you must wait until they are completely booted.
2. Any external devices attached to the head node or compute nodes such as external RAID arrays and tape drives
3. Network Attached Storage (NAS) systems (wait for the boot process to complete)
4. Head node
5. Compute nodes

Compute nodes are somewhat tolerant of the timing of starting the cluster. There is an init script on all compute nodes called *cluster-rpc-wait* that will wait up to several minutes during the boot process for the head node to become available. This actually enables you to turn the head node and compute nodes on simultaneously, and the nodes will wait for the head node. They will give up and continue the boot process after approximately 10 failed attempts at connecting to the head node.

Once the compute nodes have finished booting, your cluster should be running and ready for use. Refer to the *Cluster Software* section for information on using the cluster or *Troubleshooting* if it appears that your cluster is not running properly.

Storage Administration

Clusters and storage servers with a high quantity of hard drives will inevitably experience drive failures. Minimizing or completely avoiding downtime is entirely dependent on which preventive measures are taken. Microway strongly recommends that regularly scheduled backups are made of all vital data.

Individual Hard Drives

You should monitor the drives for signs of failure using SMART data. As root, run `fdisk -l` to determine which hard drives are in your system. SATA and SAS hard drive and SSD names will begin with `/dev/sd`. Run `smartctl` for each of those drives:

```
smartctl -a /dev/sda          (SATA/SAS Hard Drive)
```

Much of the SMART data is only useful to hard drive technicians, but an overall health assessment may be found near the top of the output. This overall-health self-assessment will read `PASSED` if the drive's health is good. Non-zero values for the `Reallocated_Sector_Ct` and `Spin_Retry_Count` SMART attributes suggest that drive failure may be imminent. It is also wise to monitor the `Temperature_Celsius` attribute to ensure the drive is properly cooled. For SSDs, the attributes `Workld_Media_Wear_Indic` or `Wear_Leveling_Count` show remaining write life. Devices reaching the end of their life should be replaced prior to failure. Please read the `smartctl` and `smartd` man pages for complete details on drive monitoring and testing with automated email alerts.

NVME drive names follow the pattern of `/dev/nvmeXnY` where X is the number of the device starting from 0 and Y is the number of the namespace starting from 1. Namespaces are similar in concept to drive partitions but operate at a hardware level making the drive show up as multiple devices. Typical usage has only 1 namespace. The first NVME drive in a system is named `/dev/nvme0n1p1`. Partitions on nvme devices get named `/dev/nvmeXnYpZ` where Z is the partition number starting from 1. For example, `/dev/nvme1n1p1` is the first partition on the first drive. NVME drives can be queried with "`smartctl -a /dev/nvme0n1`" as well but do not provide the same SMART attributes. The cli utility `nvme` can be used for additional management functionality. Reference the `nvme` man page for additional details.

Software RAID

Linux software RAIDs utilize collections of standalone hard drives to form a virtual device. The individual members of the array are still visible, so you may monitor each drive individually as described above. The health and status of the array itself may be monitored with the `mdadm` utility. Please read the `mdadm` and `mdadm.conf` man pages for detailed information on software RAID management and monitoring.

Status and configuration information for all active Linux software RAID is listed in the file `/proc/mdstat`.

ZFS

Microway provides ZFS on some configurations. ZFS is an advanced copy on write (CoW) filesystem that integrates software RAID functionality. It enables creating a single pool from 1 or more vdevs, where each vdev can be configured with the equivalent of RAID 1 (mirror), 5 (raidz1), 6 (raidz2), or triple parity (raidz3). It supports snapshots, quota/refquota, and creating multiple filesystem mount points from the same pool. For Red Hat style distributions, the ZFS on Linux repository is added to provide the required packages. OS updates may require updating the repository to a newer version at the time of updating. The utility `zpool` can be used to manage the pool and `zfs` can be used to manage the filesystems within the pool. Refer to the `zpool` man page, the `zfs` man page, and <https://zfsonlinux.org> for additional information.

RAID/NAS Systems

Hardware RAID arrays should be configured and monitored from the management interface provided with the RAID controller. Microway technicians place a link to this management interface on the desktop of the microway user (this shortcut is usually stored in `/home/microway/Desktop/`). LSI Storage Authority (LSA) is the gui management software for LSI controllers and `archttp64` is for Areca. In addition to the gui tools, cli utilities are provided. LSI uses `storcli64` and Areca uses `cli64`. Default passwords and connection information for the web guis are listed on the following page

<i>RAID Controller</i>	<i>Username</i>	<i>Default Password</i>	<i>Management Address</i>
Broadcom (LSI) LSA	root	Local root pass	<i>http://localhost:2463</i>
Areca	Admin	0000	<i>http://localhost:81/</i>

Although the RAID controller will monitor the hard drives for failures, periodic *Verify* operations should be executed on the array. Scheduled automatic RAID verifications should be configured using the RAID management utility. The verify or consistency check operation will read the data from each of the drives to ensure that everything has been properly written and is still readable. Should any sector of a hard drive fail, the verify will detect the problem and remove the faulty drive from the array. Because the verify operation may remove drives from the array, backup operations should be scheduled *before* verify operations.

The verify operation is vital because it ensures that the array will rebuild properly in the event of a drive failure. If a drive were to fail when no verify had been run recently, there is a chance that sectors on another drive had failed without being detected. In such a situation, when the RAID rebuild starts, these faulty sectors can halt the rebuild and at least some of the data may be permanently lost.

To ensure that RAID issues are appropriately handled as they arise, the **email alert settings should be configured**. The *Settings* section of the RAID management interface should be configured to email all appropriate system administrators when a failure occurs. It is vital that RAID controller warnings and errors are not ignored. When a hard drive failure is ignored, the chance of data loss greatly increases.

IPMImon™ Hardware Monitoring and Control

A Microway IPMI enabled cluster offers benefits which are unique to the industry:

- Automatic monitoring of system temperatures, voltages, and fans.
- Monitored values can be observed and power controlled from a web browser from any remote location.
- Cluster or individual nodes can be powered on and off, rebooted, and reset from a remote system.
- Out-of-bounds values are uniquely displayed.
- Physical measurements are integrated with Ganglia reporting software.
- Common operating system tasks, such as shutdown, are integrated with the user interface.
- Access to a serial console for every system in the cluster for troubleshooting or BIOS configuration
- Sends email to administrator when system measurements exceed preset limits. In extreme cases, shuts down and powers off systems. **This feature is disabled by default and must be activated by the customer. For activation instructions, see the IPMImon™ Software Configuration section.**

Cluster Software

Linux Configuration

Introduction

Microway clusters are available with many different flavors of Linux. While this manual attempts to be a generic reference for any cluster, minor details may not exactly match your cluster. If you encounter problems, please refer to your Linux distribution's documentation and support forums or contact Microway technical support.

Please do not change the host names or domain names of your head node or the compute nodes!

These names were set by Microway based on the Software Questionnaire you filled out. Changing the names will cause several of the cluster services to stop functioning correctly. You may change the appearance of the command line prompt by setting the environment variable `PS1`:

```
microway@head:~# export PS1="\u@MyNewPromptName:\w\$ "
```

```
microway@MyNewPromptName:~#
```

To make your prompt settings permanent for all users, add the export line to your `/etc/profile` file.

Environment Variables

Proper operation of your cluster requires that all necessary shell settings are correctly configured. These settings may be located in several different files, and affect either all users on the system or only a single user. Settings for all users are stored in the `/etc/profile` file and the many files in `/etc/profile.d/`. Settings for each user will be in the `.bashrc`, `.bash_profile`, or `.cshrc` files in the user's home directory (e.g., `/home/microway/`).

Any settings that were added by Microway will be found in `/etc/profile.d`. Many clusters utilize `mpi-selector` to provide multiple versions of MPI to users of the system. Other systems provide environment modules using `Lmod`. The following sections cover the 2 utilities..

MPI-Selector

The utility `mpi-selector` uses configuration files in `/var/mpi-selector/data` to enable users to select from multiple MPI libraries on the system. If your cluster uses `mpi-selector` please reference the following commands for usage:

<code>mpi-selector --help</code>	help messages
<code>mpi-selector --list</code>	available options
<code>mpi-selector --query</code>	current selection
<code>mpi-selector --set <mpi choice> --system --yes</code>	set system default(as root)
<code>mpi-selector --set <mpi choice> --yes</code>	set user selection

Lmod

If your cluster is configured with `Lmod` environment modules, Microway created `/mcms/modulefiles` as a central NFS location for all module files. Some third party optional software components such as the Intel OneAPI suite utilize their own module file directories as well. When used, Microway adds them to the `Lmod` search path on the cluster as well.

A standard set of modules for system wide use is configured in `/mcms/modulefiles/StdEnv.lua` to load a default MPI library and other selections for all users. Users may run `"module save default"` to save their own set of personal defaults once they have loaded and/or removed the modules needed for their use.

Example module commands and their purposes:

<code>module reset</code>	clears all modules
<code>module avail</code>	shows available modules
<code>module load <modulename></code>	loads module
<code>module unload <modulename></code>	unloads module
<code>module use <path></code>	add a directory with more module files, making them available for loading
<code>module save default</code>	save user defaults, used instead of system defaults
<code>module save <customname></code>	save a named set of modules
<code>module restore <customname></code>	load a named set of modules
<code>module restore system</code>	revert to system default modules

Additional documentation can be found at <https://lmod.readthedocs.io>

System Services

Linux services/daemons are started and stopped on older SysV based distributions using initialization scripts in the `/etc/init.d/` directory. Newer systemd based distributions use service files in `/usr/lib/systemd/system` but will also start scripts from `/etc/init.d` in backwards compatibility mode. `systemctl` is the primary command for interacting with systemd services including starting and stopping services.

- Systemd distros
 - RHEL/CentOS/Rocky 7+
 - OpenSUSE Leap, 13.x and newer
 - Ubuntu 15+
 - Debian 8+
- SysV distros (this list and older)
 - RHEL/CentOS 6 (upstart SysV compatibility)
 - OpenSUSE 12.x
 - Ubuntu 14.04 (upstart SysV compatibility)
 - Debian 7

You should check on-line for your distribution if not listed above or are not sure which level your system is on.

For a list of all services, run:

<code>systemctl list-units --type=service</code>	systemd
<code>chkconfig --list</code>	SysV init – Redhat/CentOS, Fedora, Suse

To enable a service (e.g., `sshd`):

<code>systemctl enable sshd.service</code>	systemd
<code>chkconfig --add sshd</code>	(RHEL/CentOS 6)
<code>update-rc.d sshd defaults</code>	(Debian 7/Ubuntu 14)

Tip: adding “--now” to a `systemctl` enable or disable command will start or start the service as well as enable or disable it

To disable a service:

<code>systemctl disable sshd.service</code>	systemd
<code>chkconfig --del sshd</code>	(RHEL/CentOS6)
<code>update-rc.d -f sshd remove</code>	(Debian 7/Ubuntu 14)

To start a service:

<code>systemctl start sshd.service</code>	systemd
<code>service sshd start</code>	init

To stop a service:

<code>systemctl stop sshd.service</code>	systemd
<code>service sshd stop</code>	init

Configuration Files

Most system configuration files will be found in /etc. Additional configs can be found in:

<code>/etc/sysconfig/</code>	(RHEL/CentOS and SUSE),
<code>/etc/default</code>	(Debian/Ubuntu)

The configuration file for X-Windows is `/etc/X11/xorg.conf` although more generic configurations may not even have an `xorg.conf` file anymore. Standard configurations are handled with automatic detection on modern distributions.

Your cluster's network has already been configured as requested in the Microway Software Questionnaire. The ethernet interface `eth0` is used by the head node and each compute node for cluster communication. The interface `eth1` on the head node acts as the gateway between your organization's network and the cluster. On the nodes, `eth1` will typically be unused. For clusters with channel bonding, both `eth0` and `eth1` will be used for cluster communication, while `eth2` on the head node will be the gateway to the outside network. If you need to change your head node's external network configuration, edit:

<code>/etc/sysconfig/network-scripts/ifcfg-eth1</code>	(RHEL/CentOS)
<code>/etc/sysconfig/networking/ifcfg-eth1</code>	(SUSE)
<code>/etc/network/interfaces</code>	(Debian/Ubuntu).

Important Linux Kernel Upgrade Note

In some special cases, Microway's software integrators replace the standard kernel shipped with your Linux distribution with a custom kernel configured specifically for your hardware. In these cases, to ensure your hardware is fully supported and stable, Microway provides a newer kernel directly from the official Linux kernel's website <http://www.kernel.org/>.

If you have a custom kernel, using your distribution's software update tool to upgrade your kernel may render your system unbootable! If you need to upgrade your kernel, please contact Microway technical support.

To determine if you have a custom kernel or a standard distribution kernel, reference the notes in the letter the accompanied your cluster. Alternatively, check the output of `'uname -r'`. If you have a custom kernel, the name should include a custom name to indicate that fact. Details of the custom kernel and the reason it was needed will be included in the letter accompanying your order.

Before shipping any cluster, Microway technicians extensively test the hardware and software on the head node and each compute node. We ensure that the kernel shipped on your cluster is stable and compatible with the hardware ordered. If your system is experiencing problems, please see the *Troubleshooting* section.

Microway Cluster Scripts

The following administration scripts are included with Microway clusters: `do.ssh`, `scom`, `scom-parallel`, `scom-nodes`, `scom-nodes-parallel`, `scomX`, `scpf` and `scpf-nodes`. These scripts utilize `ssh` or `scp` to run commands or copy files on all nodes in the cluster. Older clusters utilized `rsh` with the scripts `do.rsh`, `rcom`, `rcom-parallel`, `rcom-nodes`, `rcom-nodes-parallel`, `rcomX`, `rcpf` and `rcpf-nodes`. Newer clusters provide the `rcom*` and `rcpf*` utilities as symlinks to the `scom*` and `scpf*` versions for backwards compatibility. See the examples at the end of this section.

Each script operates on the head node and all the compute nodes or on just the compute nodes. The list of nodes, specified in `/etc/nodes`, can be overridden by setting the environment variable `NODES`, which should point to a file containing a list of nodes. For example, to instruct the scripts to run only on the first two nodes, create the file `/tmp/nodes` which contains:

```
node1
node2
```

Then export `NODES` by running:

```
export NODES=/tmp/nodes      (bash, sh)

setenv NODES /tmp/nodes      (csh, tcsh)
```

Now any Microway `scom/scpf/rcom/rcpf` script that is run in this particular shell will only execute commands on these three nodes.

Detailed descriptions of each script's function are listed below. Remember that the `rsh` versions (or symlinks) of these scripts function identically.

`do.ssh`

Simple network connectivity test. Cycles through the list of nodes attempting to login to each one. Use this script to test that `ssh` is working properly on all nodes and that network communications are functional.

`scom`

Executes a command, or set of commands, on the head node and all compute nodes. The commands are run sequentially on each system, so no commands are executed on the first node until the commands have finished on the head node. Likewise, commands are not executed on a given node until the previous node has finished the commands. `scom` accepts the parameter `-a` to prepend the hostname in each output line.

`scom-parallel`

Identical to `scom`, with the exception that the commands are run in parallel. All the systems will start running the commands at the same time, and any output from the commands will be returned in the order it is completed, not ordered by node. `scom-parallel` accepts both the `-a` parameter to prepend names and also a `-w` parameter that will wait until all commands complete prior to returning to the command prompt

`scom-nodes`

Executes a command, or set of commands, on all nodes, but not on the head node. Like `rcom`, the commands are run sequentially on each system. Thus, commands are not executed on a given node until the previous node has finished the commands.

`scom-nodes-parallel`

A compute nodes only version of `scom-parallel`.

`scomX`

Opens an `xterm` for the head node and each of the nodes and then executes a command, or set of commands, inside each `xterm`. This script is similar to `scom-parallel`, but allows you to easily see the output of each system in a small `xterm` sized appropriately to fit up to 64 nodes of output, depending on the command.

`scpf`

Copies a given file on the current system to a particular location on the head node and all compute nodes.

scpf-nodes

Copies a given file on the current system to a particular location on all compute nodes, but not to the head node.

When using these scripts, it is vital to include quotes if several parameters are being passed. Additionally, any special characters (; | < > etc.) will not be passed to the remote nodes if they are not in quotes. Any wildcard characters, variable names or additional quotes (* ? \$PS1 \$PATH " ` etc.) will need to be escaped with \ or they will be expanded before the command is passed to the remote nodes.

Examples:

Check the uptime status of each node:

```
scom-nodes uptime
```

Check dmesg for errors on all systems:

```
scom -a "dmesg | grep -i error; echo" | less
```

Sync each system's hardware clock with the Linux kernel's clock (an operation that takes a second or two):

```
scom-parallel "/sbin/hwclock --systohc --utc; rm /etc/adjtime;  
(continued) /sbin/hwclock --systohc --utc"
```

To see the importance of escaping special characters and variables, compare the following. We will assume a 4 node cluster and show example output:

```
[root@head ~]# scom "echo `hostname`"  
head  
head  
head  
head
```

```
[root@head ~]# scom "echo \`hostname\`"  
head  
node1  
node2  
node3
```

```
[root@head ~]# scom "echo $HOSTNAME"  
head  
head  
head  
head
```

```
[root@head ~]# scom "echo \${HOSTNAME}"  
head  
node1  
node2  
node3
```

Microway Cluster Management Software (MCMS™)

Introduction

MCMS is designed to enable you easily monitor and manage your cluster from any location. Critical cluster statistics are displayed in an easy to read manner. When integrated with IPMImon, MCMS provides remote monitoring and management capabilities for most hardware components. MCMS's node backup and restore functionality supports replacing failed hardware and extending your cluster with additional identically cloned systems.

Integration with IPMImon™

If your cluster was configured with IPMI monitoring hardware, MCMS will provide additional cluster statistics including system fan speeds and temperatures. The MCMS Control Panel will allow remote power on, power off and reboot as well as automated emergency reports. Clusters with IPMI offer a remote serial console session to any of the compute nodes in the event that a node does not respond via the network.

Node Backup and Restore

MCMS provides the ability to backup your compute nodes to the head node. A full metal backup is performed such that the image can be restored to a blank or used drive in a compute node. Our restore process utilizes a custom network booted environment that automatically partitions the drives, creates filesystems, extracts the image, and finally makes it bootable.

LDAP Configuration

User management is configured using OpenLDAP. The head node is configured as the LDAP server. For most situations we recommend that you use the MCMS account management web interface to manage users for your cluster. The configuration files for ldap are:

- /root/Microway/LDAP-passwords
 - Microway creates this file during setup with the 3 randomly generated passwords used for LDAP. There is a password for a read only nssproxy account (used for client authentication), a replication account (if you wish to replicate the server for backup purposes), and the ldap admin account.
- /var/www/html/mcms/admin/config.php
 - This file tells MCMS to use LDAP for authentication amongst other configuration options
- /etc/mcms/settings
 - This file contains the information for the LDAP server connection; ldaphost, root dn, cn, etc.

All other LDAP configuration files are standard to OpenLDAP. Their location and use can be found on the OpenLDAP documentation site <http://www.openldap.org/doc/>

LDAP CLI management

The command line utilities **mcms_user_add**, **mcms_user_del**, **mcms_user_mod**, **mcms_group_add**, and **mcms_group_del** can be used to manage LDAP users of the cluster. Each of the utilities has a built in help output that will be displayed with the -h option. For example, here is “mcms_user_add -h”:

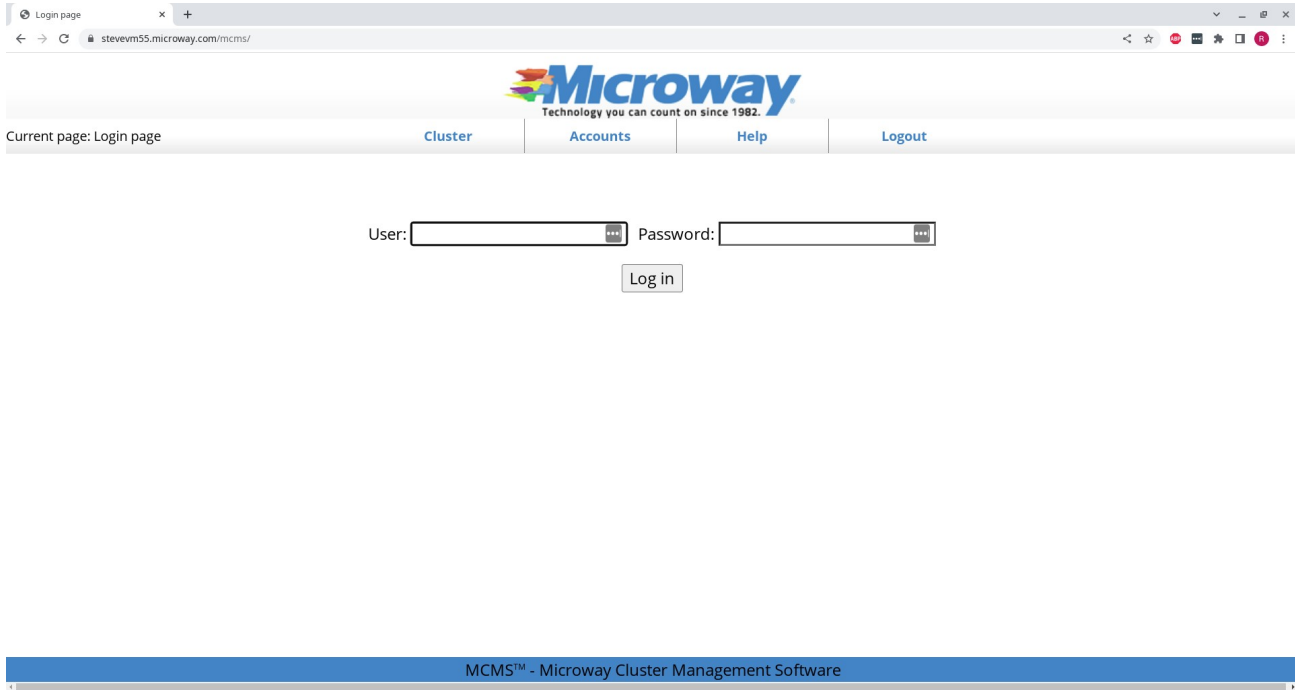
```
Usage: /usr/local/sbin/mcms_user_add [options] LOGIN

-u, --uid UID                user ID of new account (default: auto)
-g, --gid GID                group ID of the new account (default: auto)
-G, --groups GROUPS         comma separated list of additional group names
-d, --dir DIR                home directory (default: /home/LOGIN)
-s, --shell SHELL            users shell (default: /bin/bash)
-c, --comment COMMENT        GECOS field of the account, typically full name
                             defaults to LOGIN if not provided
-p, --password PASSWORD      plain text password (default: prompt for pass)
-P, --randpass               generate random password, displayed in summary
-C, --forcechange             require the user to change the pass 1st login
-q, --quiet                  do not print summary of user addition
-v, --verbose                 print additional details in summary
-h, --help                   print this help message
```

Advanced users wishing to perform actions beyond the features provided are recommended to reference the `ldif` files and `ldap` initialization script in `/usr/share/mcms`. Exercise extreme caution when working with your LDAP database manually.

Getting Started:

Open a web browser to `https://head-hostname/mcms/`, where *head-hostname* is the Internet address of the head node of your cluster (*localhost* can also be used when working from the head node directly). Once the login page loads you may enter you MCMS administrative username and password to access the management interface. The default user “microway” is configured as an MCMS admin as shipped. MCMS stores the allowed administrative users in the file `/etc/mcms/users`.



The screenshot shows a web browser window with the address bar displaying `stevevm55.microway.com/mcms/`. The page features the Microway logo at the top center, with the tagline "Technology you can count on since 1982." Below the logo, a navigation bar includes links for "Cluster", "Accounts", "Help", and "Logout". On the left side of this bar, it says "Current page: Login page". The main content area contains a login form with two input fields: "User:" and "Password:", each followed by a password toggle icon. A "Log in" button is positioned below the password field. At the bottom of the page, a blue footer bar displays the text "MCMS™ - Microway Cluster Management Software".

Cluster Control

The *Cluster Control* page presents IPMImon hardware monitoring data and simplifies remote administration of the cluster. After the MCMS login process is complete, the *Cluster Control* page will be shown. This initial page shows controls for cluster administration and cluster shutdown/restart. From here, the administrator can copy files between portions of the cluster and execute commands across the cluster. You may reboot or power off select nodes or the entire cluster. At the bottom of the page, fan and temperature data will be displayed if your cluster is equipped with IPMImon.

Cluster Control

Current page: Cluster Control

Cluster Accounts Help Logout

Copy file from head to node(s)

From: To path: All Copy

Perform a command on node(s) as root

Command: All Execute

Shutdown node(s): None Shutdown

Reboot node(s): None Reboot

NodeWatch™

Nodes	Fans	Volts	Temps	OS			
Name	FAN3	FAN5	FANA	FANB	Status	Status	Gstat
head	1900	1000	1000	1700			
node1	1800	900	800	1500			
node2	1800	900	900	1600			
node3	1700	800	800	1400			
node4	1700	800	800	1500			

Power Off Power On Reset Console

MCMS™ - Microway Cluster Management Software

A navigation bar is available across the top of this page.

- Cluster menu provides access to:
 - Cluster Control
 - Node Backup
 - Node Restore
- Accounts menu allows for managing:
 - Users
 - Groups
- Help menu provides:
 - A PDF copy of this manual
 - Version information
- Logout

To copy files from the head node to the compute nodes, type the full path of the source file (on the head) in the *From* box and then the full path to the destination file (on each compute node) in the *To* box. If you only want to copy the file to one particular node, select it from the pull-down menu. Then select *Copy*.

Execution of commands across the cluster is the same. Simply enter the command in the *Command* box and then select what systems the command should be executed on. Remember that selecting "All" will execute on the head node and the compute nodes. As the command completes, you will see the output from each system.

To shutdown or reboot a system or group of systems, select them from the pull-down menu and then select *Shutdown* or *Reboot*. **This will immediately begin the shutdown process, so first ensure that your**

users are not running any jobs and have saved all their work. If you selected shutdown, the systems will power off cleanly via software. Shutting down or rebooting your head node while the nodes are running is generally not recommended as many of the nodes' services rely on the head node. If the head node is rebooted, the simplest way to ensure everything is working properly on each node is to reboot all the nodes.

On a cluster equipped with IPMI, the *Cluster Control* page will list the fan and temperature status of each node. The Fans, Volts, and Temperature column groups can be expanded or collapsed to display either individual values or aggregate health status. When the IPMImon daemon is unable to communicate with a node, or if a measured value is out of range, this will be indicated on the display with unique colors. Orange is used to denote data that is out of range (the system is too warm, or a fan is spinning too slowly). Red is used to denote measurements that will cause IPMImon to consider shutting down the affected system (the system is dangerously hot, or a fan has stopped spinning). It is normal for nodes to take a moment to report their first set of data after restarting the system. If problems persist, refer to the *Troubleshooting* section.

Select a node to use the buttons at the bottom for powering off/on the selected node **Please keep in mind that these buttons are equivalent to pressing the physical switch on the front of the system. Under most circumstances, you should use normal operating system commands to shutdown or reboot a node. You may use the *Shutdown* and *Reboot* menus and buttons near the top of the *Cluster Control* page, which use the operating system to shutdown or reboot the system.**

There is also a Console dropdown menu which lets you select from available console choices. The VNC console option starts a VNC session on the selected node and provides it through the web gui. Serial console support is also available for essentially all IPMI enabled clusters. It provides low level access including access to the POST and BIOS screens during system reboots. Even when a node is unavailable over the gigabit ethernet and InfiniBand networks, the serial console can still be used in most cases. Microway also has an experimental iKVM passthrough feature enabling access to the IPMI KVM providing a full graphical remote access console. This feature is not available on all platforms and will be enabled only when supported.

Current page: Cluster Control

[Cluster](#) [Accounts](#) [Help](#) [Logout](#)

Copy file from head to node(s)
 From: To path: All

Perform a command on node(s) as root
 Command: All

Shutdown node(s):
 Reboot node(s):

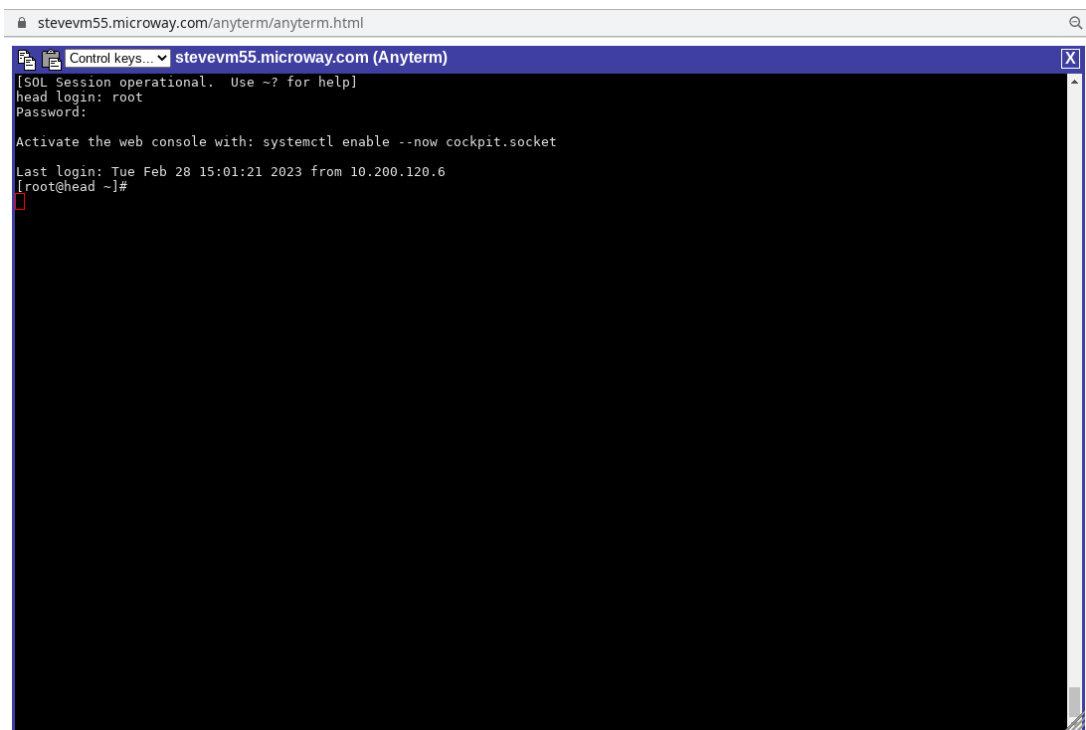
NodeWatch™

Nodes		Fans				Volts	Temps	OS
Name	FAN3	FAN5	FANA	FANB	Status	Status	Gstat	
✓ head	1900	1000	1000	1700	●	●	●	
node1	1800	900	800	1500	●	●	●	
node2	1800	900	900	1600	●	●	●	
node3	1700	800	800	1400	●	●	●	
node4	1700	800	800	1500	●	●	●	

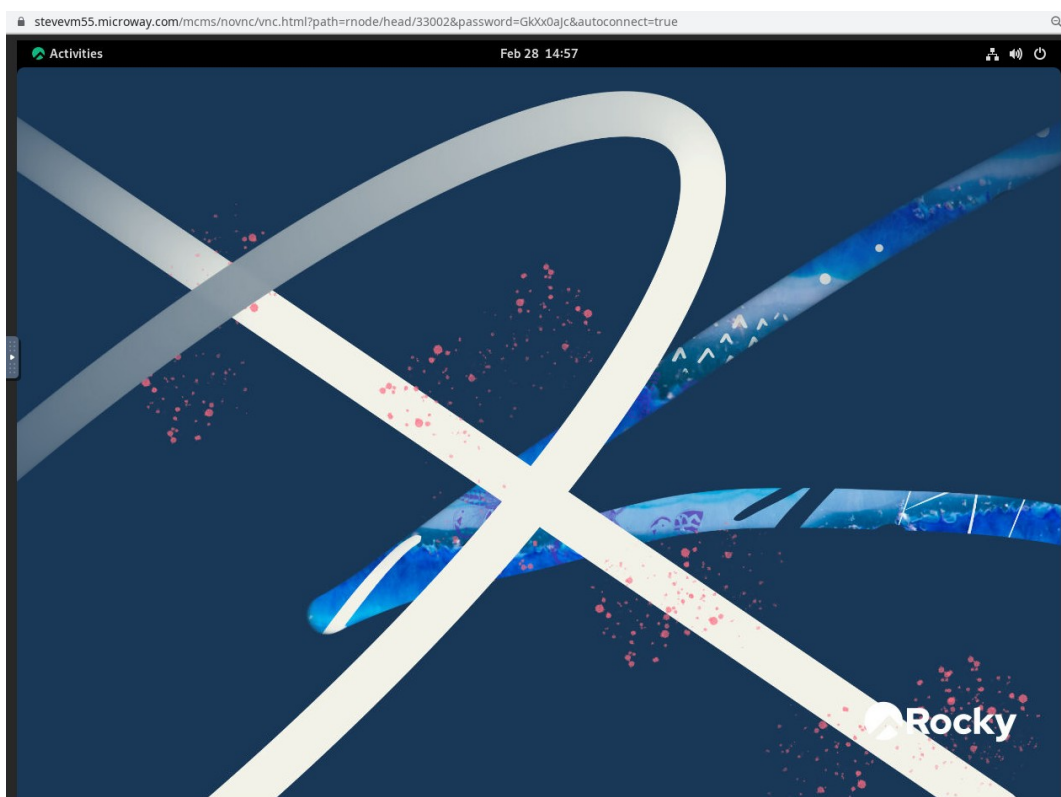
Serial
Console
iKVM
VNC

MCMS™ - Microway Cluster Management Software

Console dropdown menu



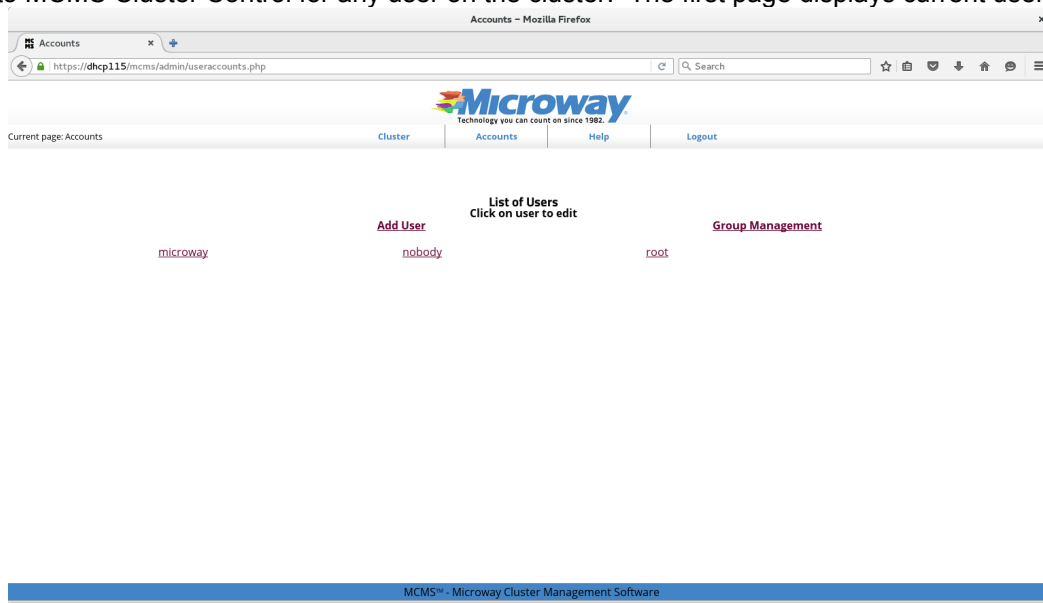
Serial console window



VNC Session

Account Administration

Selecting *Account Admin* from the *Cluster Control* menu opens the user account administration page. From this page, you can add, remove and modify users across the cluster. Additionally, you may grant or revoke access to MCMS Cluster Control for any user on the cluster. The first page displays current users:



To add a new cluster user, select *Add User* and enter all pertinent information:

To modify an existing user, select their name and change the information:

Account being edited: **microway** [Delete](#)

Full name:

Password: Repeat:

Shell:

User must change password at next login? ☐

User is an MCMS administrator? ☒

Groups user is a member of:

[Update](#) [Reset](#)

MCMS™ - Microway Cluster Management Software

Any user that has a check next to “User is an MCMS administrator?” will be able to enter the MCMS Cluster Control page and make changes to any section. **You should not check “User is an MCMS administrator?” when adding normal users!**

Node Backup

MCMS allows you to easily create a backup image of any of your nodes and store the image on your head node. This image can then be used to re-image any of the nodes in your cluster. This feature will greatly simplify any major changes that you would like to make to the nodes and will be useful when adding nodes for any future cluster expansions. You will also find it helpful should a node need it's system hard drive replaced. To create a backup, select *Node Backup* from the menu.

The destination system must have the same or larger storage device as the source for an image to be installed on it. Manual modifications to the node image can be performed to adjust sizing or target devices if needed. Contact Microway support for assistance with these topics if required.

Node backup and restore provides support for both legacy and EFI boot modes. You must boot the system in the mode matching the backup image for successful restoration.

MCMS Node backup

Please note that your current backup will overwrite the prior one.

Please select node to be backed up:

Please choose backup slot:

Optional comment field for this backup slot:

Modify config options if necessary:

Option explanations:

HEAD-IP of head node

DEVICES-hard drives in node to be backed up

ETHMODS-network modules to be loaded the restoration

NODERESTORE-nfs mountpoint on head to restore from

Enter any directories that you do not want backed up. Enter one directory per line.
NFS mounted directories will automatically be excluded already.
Example:
/scratch/tempdata
/data

[Backup now](#) [Reset](#)

MCMS™ - Microway Cluster Management Software

MCMS node backup page

Although there are many settings on this page, they should be preset correctly - Microway technicians have customized the page for your cluster. The default config options are set in `/etc/mcms/nodebackup-defaults`

First, you need to select which node should be backed up. Choose a node with the configuration you want to keep that is known good. **If you are performing a node backup/restore because a node experienced a hard drive failure, do not backup the node with the bad hard drive!**

Next, select which backup *slot* you would like to use to store the backup. Write a short description of the backup in the comment field to identify it.

Because Microway technicians configured the backup for your cluster in the file `/etc/mcms/nodebackup-defaults`, you should be able to skip the configuration options. Should you need to change them, their descriptions are listed below.

HEAD

Should match the IP address of the head node.

NODERESTORE

Should not be modified unless the MCMS backup directory was moved. The default directory is `/mcms`.

DEVICES

Use AUTO to allow MCMS to automatically detect used drives for the backup.

For manual control, specify the list of hard drive devices in the nodes, separated by spaces. Drives will be `"/dev/sda /dev/sdb /dev/nvme0n1"` etc. To determine what drives are in your nodes, run `df` and read the left-most column. For software RAID devices, specify all array members listed in `/proc/mdstat` (e.g., `"/dev/sda /dev/sdb /dev/sdc"`). Additionally, any node with a software RAID must have its configuration in `/etc/mdadm.conf`. Examples include:

```
ARRAY /dev/md1 uuid=b23f3c6d:aec43a9f:fd65db85:369432df
```

```
ARRAY /dev/md3 level=raid5 num-devices=3 devices=/dev/sda3,/dev/sdb3,/dev/sdc3
```

ETHMODS

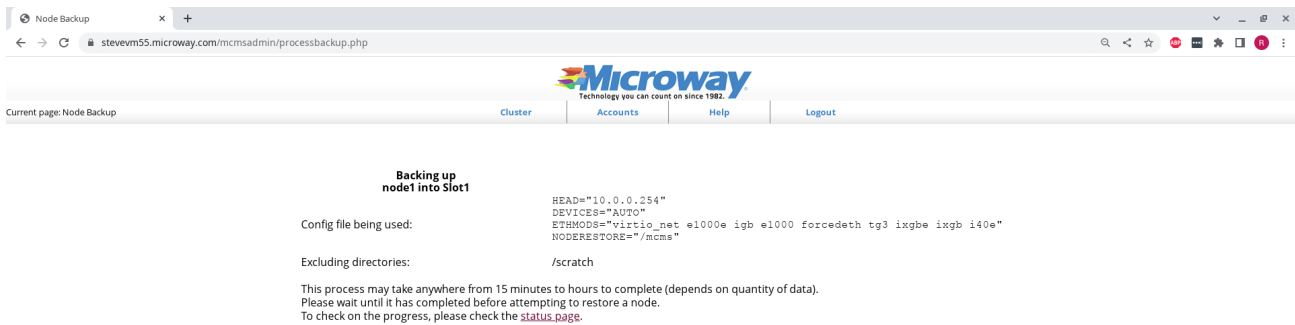
Ethernet kernel module list for network drivers. The modules will be loaded in the order listed. Any devices that are not present will be ignored. This allows using a wide range of drivers in the restore image for increased compatibility. To determine what network card is in your nodes, run `lsmod` on a node. Commonly used modules are `e1000`, `e1000e`, `ixgb`, `ixgbe`, `tg3`, and `forcedeth`.

NETDEVS (optional)

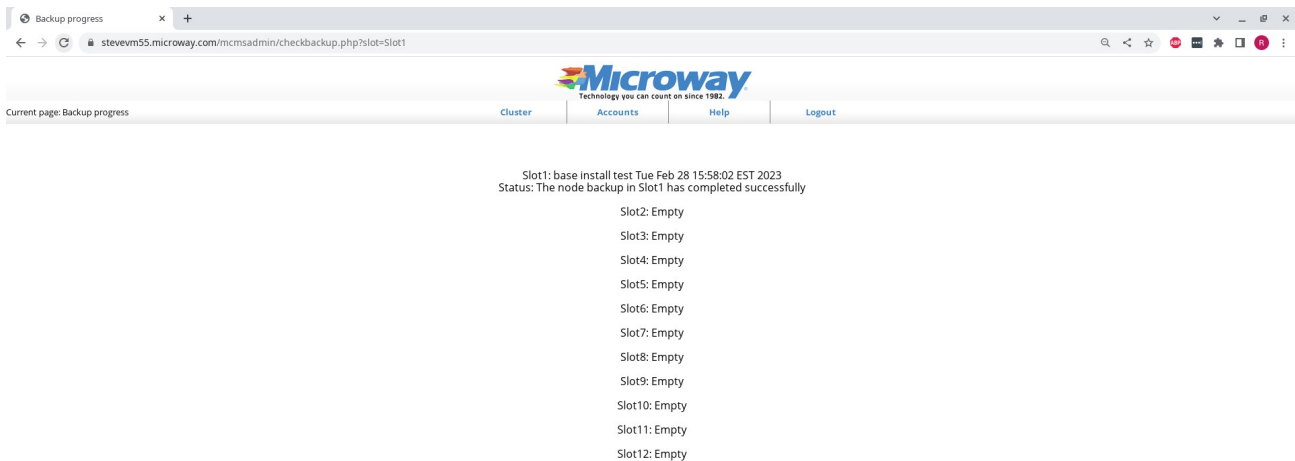
Noderestore defaults to using `eth0` for restoration. If you are using a different ethernet device, specify it here.

Finally, decide if you would like to exclude any file systems from the backup. NFS, lustre, BeeGFS, panfs, snap, and gvfs filesystems/directories are automatically excluded. If there are other network file systems, ephemeral filesystems, or scratch directories which should be excluded they are entered here. The example graphic shows the directory `/scratch` being manually excluded. Network file systems will be mounted by the nodes after they are restored, so those directories should not be backed up. To see what a node currently has mounted, run: `df -h` or `mount`.

Once you have finished configuring the backup, select *Backup now*. A summary page will be presented. If any errors occurred starting the backup, they will be displayed on the page. While the backup is running, you can view an automatically updating status by selecting *status page*. When the backup has finished, or if it encounters an error, the status page will inform you.



Backup summary page



Backup status page

Once complete, the backup image is ready to be used for a restore. You can check the sizes of the backup images on your head node by looking in the `backups` directory within the directory specified by the `NODERESTORE` backup configuration option. The default directory is `/mcms/backups`.

Node Restore

By default, the node image is set to the most recent backup slot. If you would like to restore a node using an older or different node image, select *Node Restore* from the menu on the left hand side of the page. Then choose the desired backup slot from the pull-down menu.

To restore node backup images, nodes are PXE booted from the network and will then execute an automated system restore script. When the script completes, it will prompt you for a node number. This allows you to backup any node and restore it to any other node number.

On any node that has a new, empty hard drive, the network PXE boot process will be completely automatic. However, any node that already has Linux installed will attempt to boot from the drive instead of the network. You can erase Linux from a node's hard drive and force a network boot, however, **this command will erase the partition table making all data unavailable on the hard drive:**

```
dd if=/dev/zero of=/dev/sda bs=1M count=20
```

In addition to the partition table at the beginning of the disk, GPT partitioned disks have a backup copy of the table at the end of the disk. If you need to wipe a GPT partition table completely, use the utility `sgdisk`. The below syntax will wipe the GPT tables at both the beginning and end of the drive `sda`:

```
sgdisk --zap-all /dev/sda
```

IPMI enabled systems can utilize the BMC to select network boot as the next choice. Running the following will set node5 to boot to the network in legacy mode on the next power on or reboot:

```
ipmitool -I lanplus -H node5-ipmi -U ADMIN -P yourPass chassis bootdev pxe
```

If the system is in EFI mode, the additional parameter `options=efiboot` should be added to the end of the `ipmitool` command to network boot in EFI mode.

IPMI Serial over LAN (SOL) functionality can also be used to watch the restore process and enter a node number at the end. The MCMS web gui provides a SOL interface. From that interface, if PXE booting in legacy mode you will be presented with a PXELinux boot menu. Select the "MCMS automatic node restoration (serial console mode)" option to have the restore process run on the serial console instead of the local VGA console. You can edit `/tftpboot/pxelinux.cfg/default` to change the default legacy PXE boot option. For EFI mode, a Grub boot menu will be presented where you can select "MCMS Noderestore Serial Console COM2" to use the serial console for an EFI network boot.

You can also force a network boot from the serial console or by attaching a keyboard and monitor to the node and repeatedly pressing <Esc>, <F8>, <F9>, <F10>, <F11> and <F12> immediately after the system powers on. On any modern system, one of these keys should trigger a network boot or boot menu where you can select network. <F12> is the most common key for booting directly to the network.

When the node restore has completed and asks for a node number, simply enter the last number of the node's IP address. For a default cluster configuration, node9's IP address would be 10.0.0.9. When restoring node9, you would enter 9. The system will automatically reboot after the node number has been entered.

If a node does not properly set its node number during the restore, once up you can manually set it:

```
/usr/sbin/change_ip.sh #
```

where # is the last number of the node's IP address. Reboot the node for the changes to take effect.

Other User Administration Issues

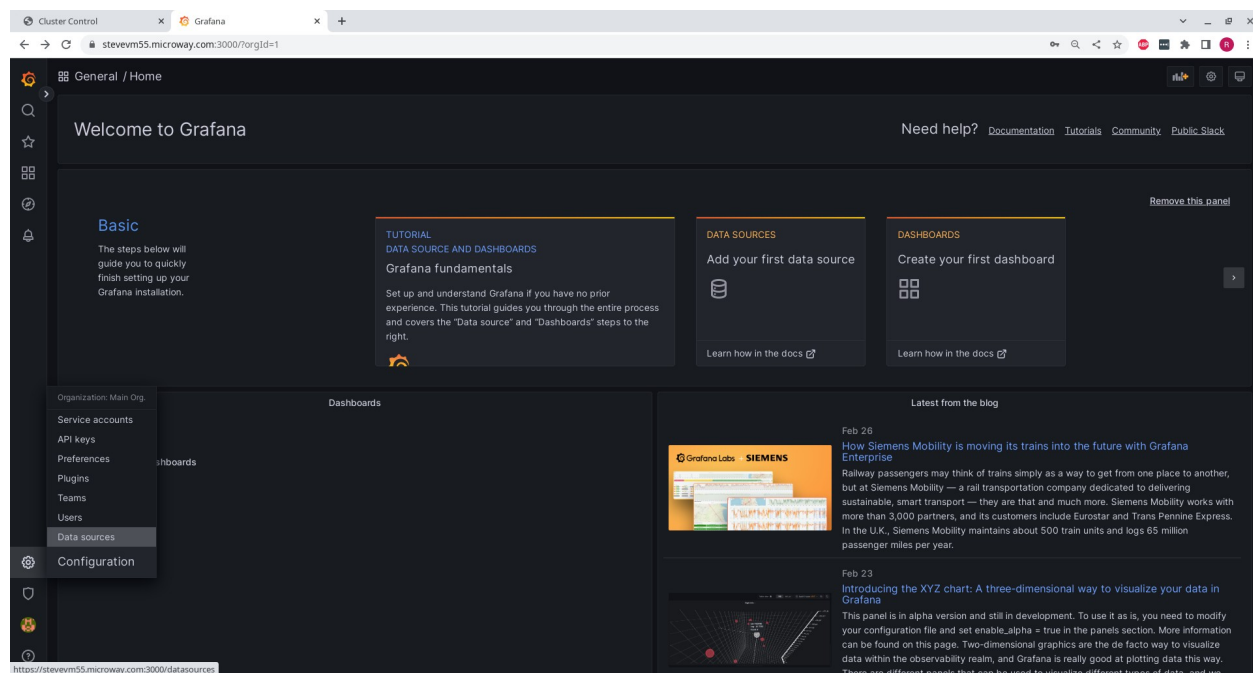
MCMS Administrative User Configuration

Access control for command execution privileges uses the operating system's user account database. Existing users who can log into the head node can be granted authorization to log in with the same username/password combination at the MCMS Cluster Control user login prompt. The user “**microway**” with password “**microway**” has already been configured as an MCMS administrator. We suggest that you change this upon receipt of the cluster - see the *Account Administration* section above. If you accidentally delete all MCMS administrators, you can simply edit `/etc/mcms/users` on the head node. This file is a list, one per line, of existing Linux user accounts that are granted MCMS privileges.

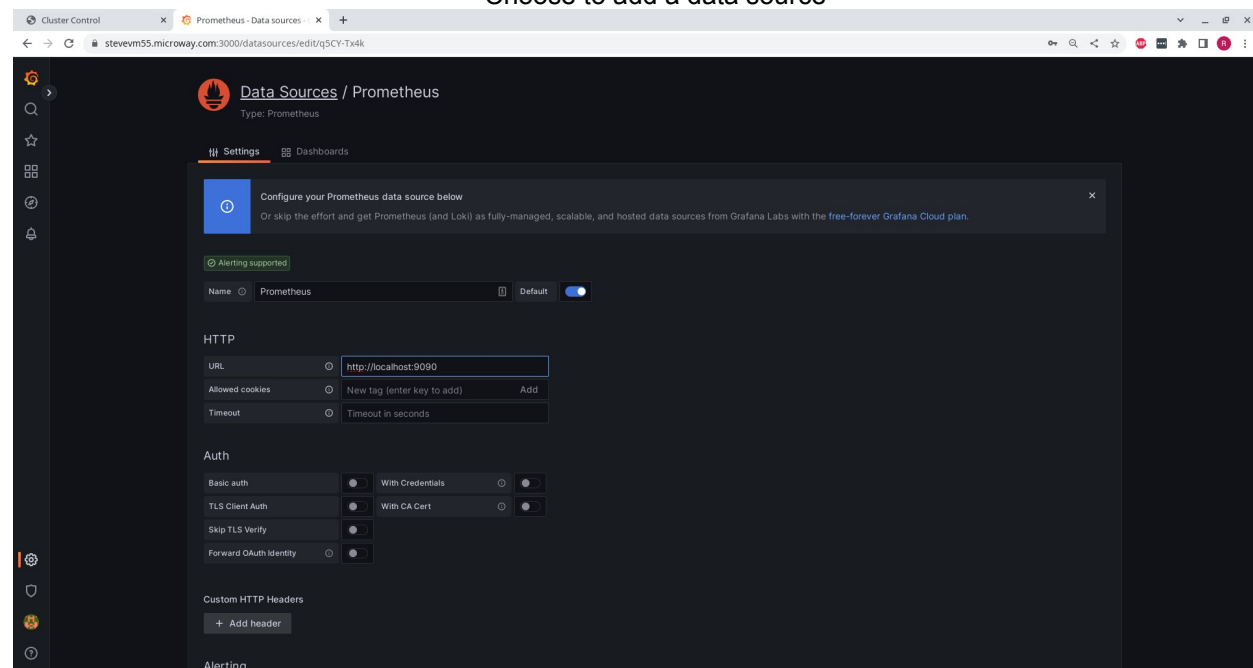
Grafana

Microway includes Grafana for viewing metrics about your cluster. Grafana is a multi-platform open source analytics and interactive visualization web application. It provides charts, graphs, and alerts for the web when connected to supported data sources.

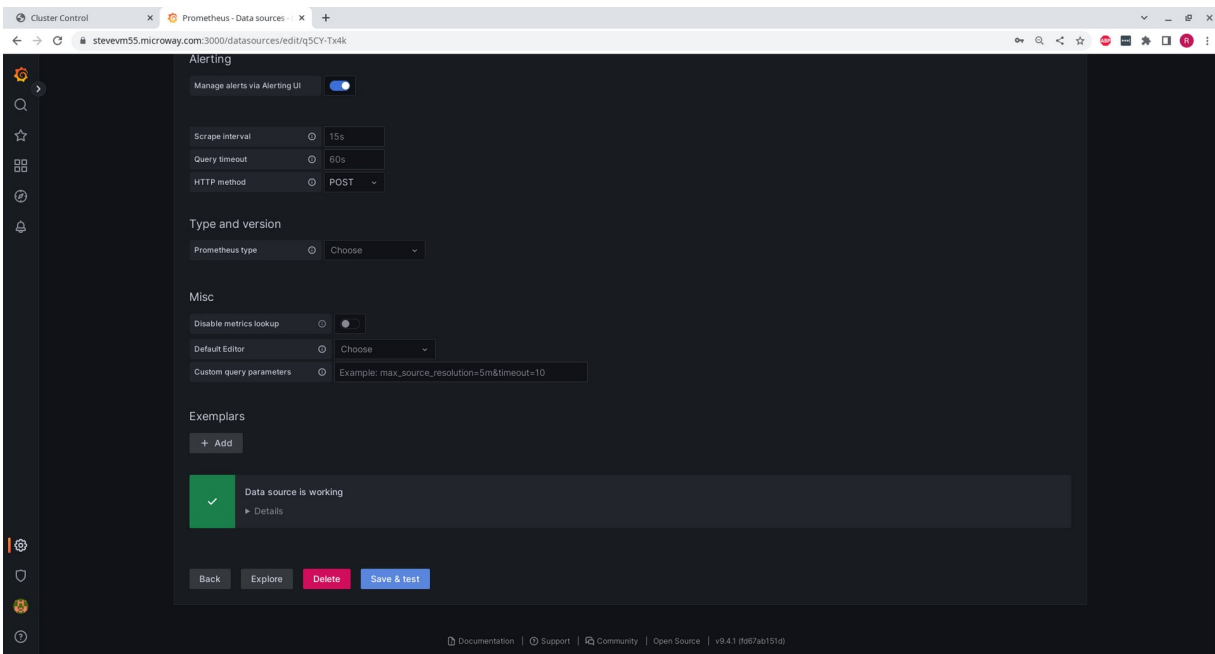
The prometheus database and data gathering service is run on the head node of the cluster. Each cluster system also runs the prometheus node exporter service. That service monitors system metrics such as system load, disk space used, memory usage, CPU usage, and network usage. The information is gathered from each system by the prometheus service. The data is then made available as a “Data Source” for Grafana. The data source is added to Grafana as follows:



Choose to add a data source

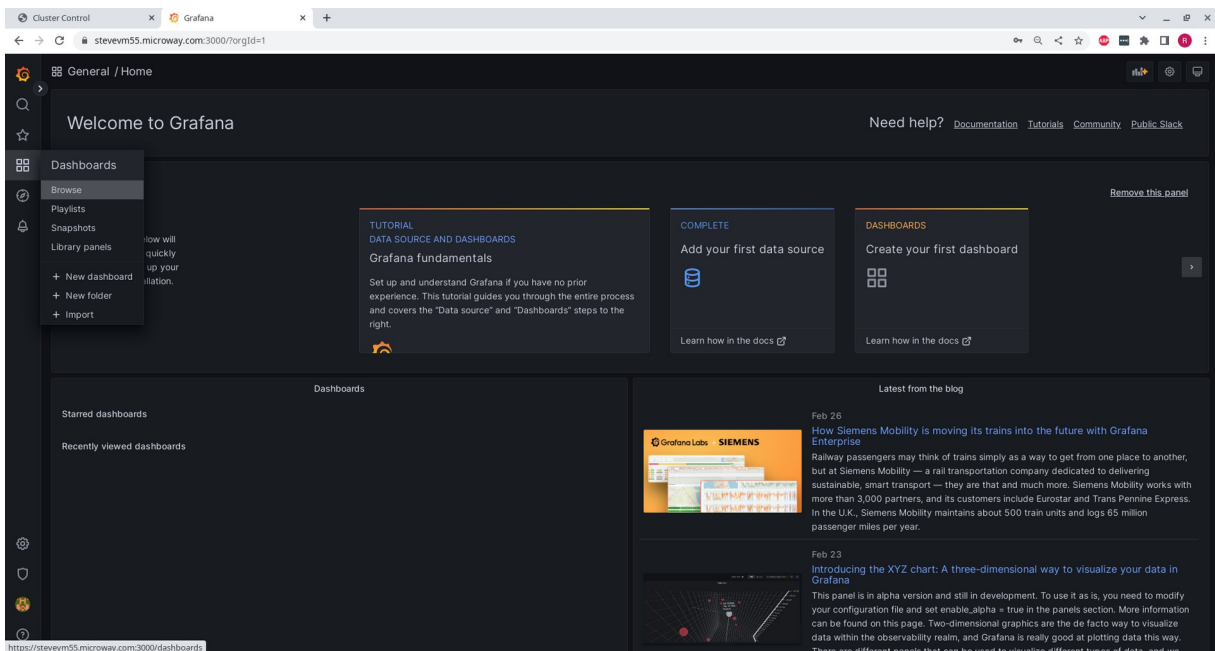


Enter the URL `http://localhost:9090` for the default Prometheus service settings

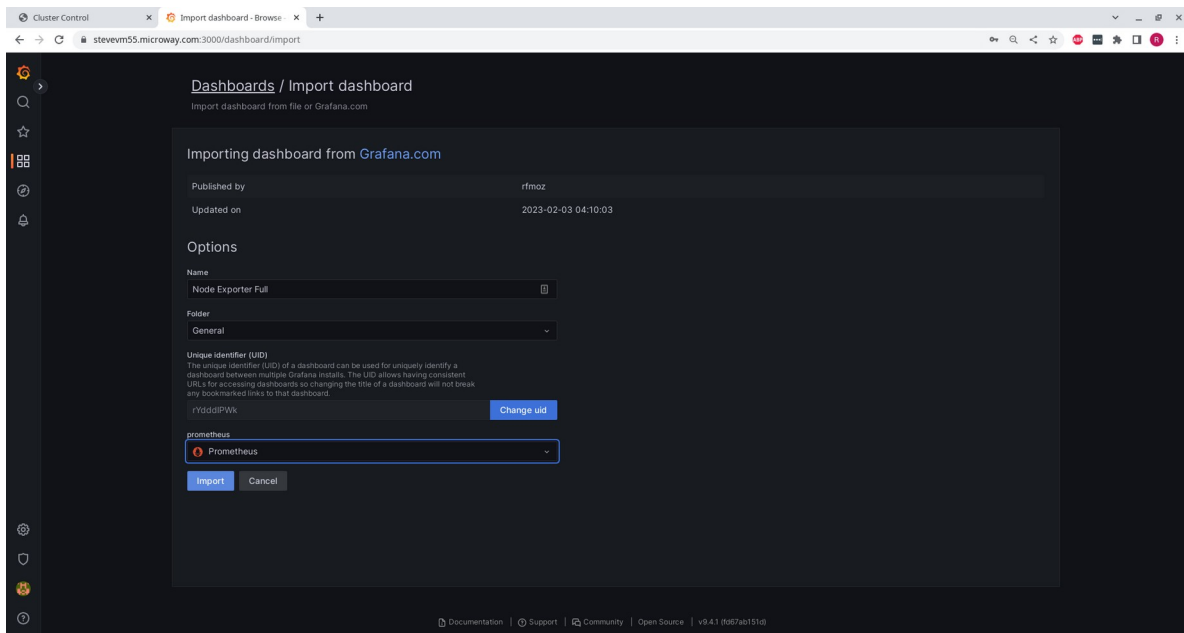


Save and test the data source

Once the data source is available, dashboards based on it can be added. Open the dashboard browse page and you can choose to import a dashboard from the “New” dropdown menu. Grafana provides a library of dashboards at <https://grafana.com/grafana/dashboards/>. To import dashboards from the repository, enter the dashboard ID in the dashboard import popup window. The node exporter dashboard ID is 1860 and is imported in an example image. When imported, you must select the data source it should use which is normally Prometheus unless you have configured additional sources yourself.

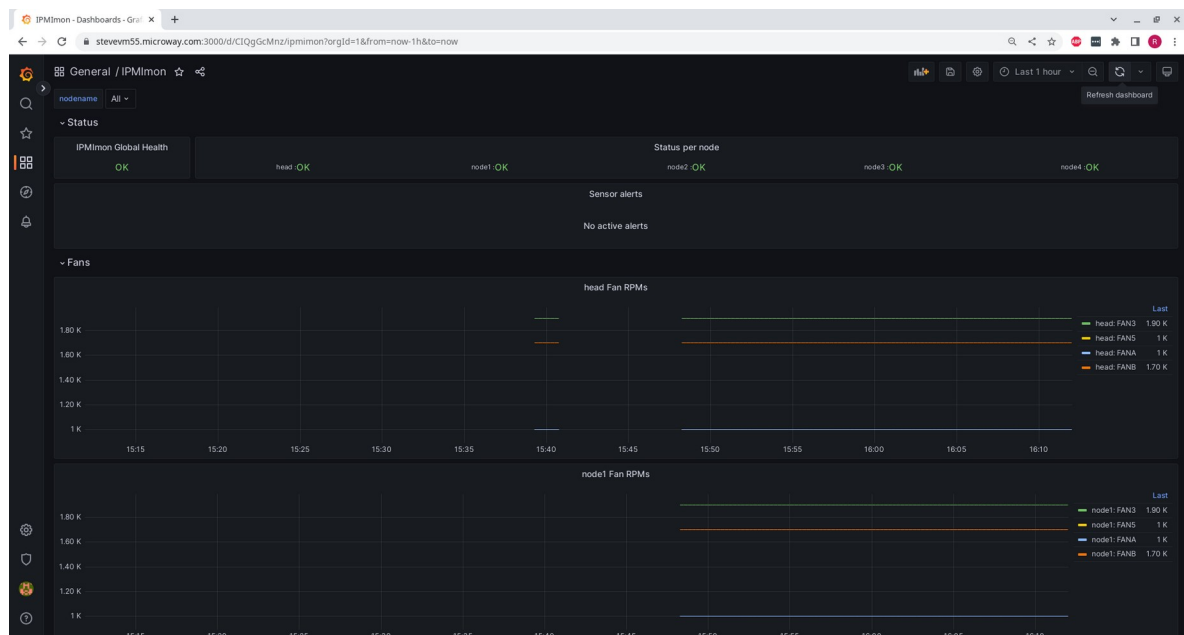


Open the dashboard management browse page



Set the data source

Microway also provides a dashboard for the data collected by IPMIImon. See the IPMIImon manual section regarding ipmimon-exporter for details on the export service. The dashboard is provided as an importable json file on the head node in `/usr/share/mcms` named `IPMIImon-<id #>.json` where `id #` references the current release number. As of this document, the id is 1657314442526.



Example data from a small cluster

We encourage you to explore the Grafana dashboard repository for useful dashboards to either use directly or to use as examples to create your own customized dashboards from.

MCMS Installation

All Microway clusters have MCMS installed before shipping. If you have reinstalled Linux on your cluster, you will first need to install mcms-network. You do not need to reinstall this package on a cluster configured by Microway! You will only need to reinstall this package to upgrade from an older MCMS version, or if you have re-installed the operating system for your cluster.

Microway provides package repositories for Red Hat/Rocky 7, 8, and 9 as well as Ubuntu 20 and 22. All MCMS repositories utilize GPG signing. Our public key is available at

<https://repo.microway.com/pub/microwayhpc-v4/RPM-GPG-KEY-MicrowayHPC>

For Red Hat style distributions, download the GPG key and import it with “rpm –import <filename>”. Rpm files installing our repository are available for the supported versions as follows:

- Red Hat / CentOS 7
 - https://repo.microway.com/pub/microwayhpc-v4/redhat/7/x86_64/microwayhpc-release-1.2-1.el7.noarch.rpm
- Red Hat / Rocky 8
 - https://repo.microway.com/pub/microwayhpc-v4/redhat/8/x86_64/microwayhpc-release-1.2-1.el8.noarch.rpm
- Red Hat / Rocky 9
 - https://repo.microway.com/pub/microwayhpc-v4/redhat/9/x86_64/microwayhpc-release-1.2-1.el9.noarch.rpm

Microway's Ubuntu repositories use the same GPG key as the RPM repositories above. Once downloaded, install it as `/etc/apt/trusted.gpg.d/microwayHPC.asc`. Both Ubuntu 20.04 and 22.04 are supported. Create `/etc/apt/sources.list.d/microwayHPC.list` with the correct contents for your Ubuntu version as follows:

- Ubuntu 20.04
 - `deb [arch=amd64] https://repo.microway.com/pub/microwayhpc-v4/ubuntu20 stable main`
- Ubuntu 22.04
 - `deb [arch=amd64] https://repo.microway.com/pub/microwayhpc-v4/ubuntu22 stable main`

Microway also provides the source the packages are built from. The source is currently managed on our internal git infrastructure and it not available on a public site. Please contact Microway support to receive a copy if desired. Build instructions can be provided at that time.

All nodes require mcms-network to be installed and configured prior to beginning the remainder of the MCMS installation process. Once completed, the compute nodes need the mcms-client and mcms-client-ldap (if using LDAP auth) packages installed. The head node needs: mcms-head (meta package includes: mcms-cert-auth, mcms-noderestore, and mcms-webgui), ipmimon, and mcms-ldap(if using LDAP auth).

MCMS-network Installation and configuration

All Microway clusters have mcms-network installed before shipping. If you have reinstalled Linux on your cluster, installing and configuring mcms-network will configure the host names, network interfaces and remote login daemons. You do not need to reinstall this package on a cluster configured by Microway!

The latest version is available in our repositories listed above. The filename is in the format `mcms-network-x.y.z-rel.(rpm|deb)`

- Install the package for your distribution using yum or apt once the repository has been configured
- Edit **/etc/mcms/network** to provide the network settings for the system. These are the available settings:
 - `head` – The hostname for the head node of the cluster. The default name is “head”.
 - `nodes` – The hostnames for the nodes of the cluster. The default name is “node”, with the final portion of each node's IP address appended. This makes the default node names node2, node3, etc. You do not need to set this or head if you are using `UsesNodefile`.
 - `nodedigitplaces` – number of places for the node number (ie node2 vs node02)

- **NetMask** – The network mask for the head node and compute nodes. The default is 255.255.254.0, which is sufficient for any cluster up to 254 nodes when using the same subnet for both the primary ethernet connection and IPMI
 - **Net** – The IP address range for the cluster. The default is “10.0.0.”, which makes the full IP addresses 10.0.0.x.
 - **HeadIP** – The IP address for the head node. Defaults to 1
 - **NodeStartIP** – The IP addresses of the cluster nodes will start at this number. The default is 2, which makes the full IP addresses x.x.x.2, x.x.x.3, etc.
 - **NumNodes** – The total number of nodes in the cluster, including the head node.
 - **Domain** – The domain name of the cluster. The default is “cl.microway.com”, which makes the full names head.cl.microway.com, node2.cl.microway.com, node3.cl.microway.com, etc.
 - **Device** – The ethernet device to be used for cluster communication. The default is eth0.
 - **UsesNodefile** – The path to a plaintext file which lists the desired names for nodes. If you wish to use node names which do not conform with Microway's default scheme, write them in a separate text file with one name per line. You must specify the full path to the file.
 - **CreateIPMInames** – Should entries in /etc/hosts be made with IP addresses and names for IPMI devices for management?
 - **IPMINet** – Network used for IPMI. This network should be reachable (same subnet) as the cluster network when using a shared network for both services
 - **CreateAltNames** – If the cluster has InfiniBand or other high speed network, set this to yes to create entries in /etc/hosts for it and also to create ifcfg-ib0 (or similar config for different distros) files to set IP addresses for the devices
 - **AltNet** – The IP network used for InfiniBand/AltNet TCP/IP. This should be a different subnet than the cluster network and IPMI network.
 - **AltDevice** – The network device used by InfiniBand. Defaults to ib0
 - **AltSuffix** – hostname suffix for the AltNet. Defaults to -ib.
 - **Mode** – This parameter controls the mode of configuration. Mode=head will enable NAT to the nodes to share the external network connection within the private cluster network. If set to Mode=compute, the nameserver and gateway will be set to the head node. Additional parameters can be modified to change this behavior.
 - **EnableGateway** – used when in Mode=compute, set the gateway to the head node if enabled
 - **EnableDnsmasq** – used in Mode=head. Configure and enable the dnsmasq service if enabled
 - **EnableMasquerade** – used in Mode=head with BasicFirewall=yes, configure and enable IP Masquerading/NAT to provide outside network access to the compute nodes.
 - **NatDevice** – if EnableMasquerade is set to yes in Mode=head, NatDevice sets the external device to share.
 - **BasicFirewall** – Used with Mode=head only. Creates a basic firewall setup with the cluster and AltNet (if used) wide open, and the LAN/WAN connection (NatDevice) allowing ports 22(ssh), 80(http), and 443(https) only.
- Run the initialization script: /usr/share/mcms/mcms-network-init.sh
 - Change the node's hostname and address: change_ip.sh 1 (use the last octet of the desired IP)
 - Reboot the head node and all compute nodes.
 - To configure the required password-less ssh, you will need to configure the ssh keys for root once this completes. On the head node, run the following command for each compute node you are configuring

ssh-copy-id node1

The sample mcms-network /etc/mcms/network file on the following pages has been configured for a cluster with 17 nodes (head + 16 compute). The head node will be head.cluster.myuniversity.edu with IP address 10.0.0.254. The nodes will be nodeX.cluster.myuniversity.edu with IP addresses 10.0.0.X (1-16). The head node and compute nodes will use their first ethernet port, eth0, for all communication. IPMI and IB names are enabled, and the config file is set to run in head mode. The same config file with Mode changed to compute could be used on the compute nodes. The script will not alter the configurations of the other ethernet ports. The head node and compute nodes will use ssh for internal communication.

```

#This config file controls the network settings to be applied.
#If you wish to use a custom list of hostnames, please specify the filename
#in UsesNodefile. If UsesNodefile is set, the head and nodes variables
#are ignored.

head="head"
nodes="node"
#nodedigitplaces=1 means node2, =2 means node02, =3 means node002
nodedigitplaces=1
NetMask="255.255.254.0"
Net="10.0.0."
HeadIP=254
NodeStartIP=1
#NumNodes is the total number of systems, including head
NumNodes=16
Domain="cluster.myuniversity.edu"
Device="eth0"
UsesNodefile=""

#Should ipmi names/ips be added to /etc/hosts?
CreateIPMIInames="yes"
#CreateIPMIInames="no"
IPMINet="10.0.1."

#Should Infiniband names/ips be added to /etc/hosts and ifcfg files be
created?
CreateAltNames="yes"
#CreateAltNames="no"
AltNet="10.1.0."
AltNetmask="255.255.255.0"
AltDevice="ib0"
AltSuffix="-ib"

Mode="head"

```



```
#Mode="compute"

#if you are in Mode=compute, set the gateway to the head
EnableGateway=yes
#EnableGateway=no

#if you are in Mode=head, enable dnsmasq
EnableDnsmasq=yes
#EnableDnsmasq=no

#if you are in Mode=head, enable iptables MASQUERADE for node's NAT
EnableMasquerade=yes
#EnableMasquerade=no

#Which devices provides the outside ethernet connection in Mode=head
#when EnableMasquerade is set to yes?
NatDevice=eth1


#if Mode=head and EnableMasquerade=yes and this is set yes, then set up basic
firewall rules
#(allow all lo, eth0, ib0 [if CreateAltNames=yes], ssh, http, https to eth1)
BasicFirewall=yes
```

MCMS head and client package installation

Once the `mcms-network` configuration and initialization is completed you can proceed to installing the remaining packages. The compute nodes need the `mcms-client` and `mcms-client-ldap` (if using LDAP auth) packages installed. The head node needs: `mcms-head` (meta package includes: `mcms-cert-auth`, `mcms-noderestore`, and `mcms-webgui`), `ipmimon`, and `mcms-ldap` (if using LDAP auth). The packages require the removal of the `gnome-initial-setup` package to enable MCMS to handle authentication instead of the built in mechanisms. The apt package manager on Ubuntu does it automatically, but you need the `--allow-erasing` parameter when using yum.

- Red Hat
 - Head: `yum install --allow-erasing mcms-head mcms-ldap ipmimon`
 - Compute: `yum install --allow-erasing mcms-client mcms-client-ldap`
- Ubuntu
 - Head: `apt install mcms-head mcms-ldap ipmimon`
 - Compute: `apt install mcms-client mcms-client-ldap`

Once the packages are installed, you must generate SSL certificates with the utility `mcms-certgen.sh`:

edit `/etc/mcms/sslspecs` to define the values used for your certificates.

Example:

```
SSLSPECS="/C=US/O=Example
Institution/ST=Massachusetts/L=Plymouth/CN=head.cl.institution.org/
subjectAltName=DNS.1=head.cl.institution.org,DNS.2=head,DNS.3=localhost"
mcms-certgen.sh -r                (create root CA)
mcms-certgen.sh -a                (create apache cert)
mcms-certgen.sh -l                (create LDAP cert)
mcms-certgen.sh -n                (distribute public root CA cert to nodes)
mcms-certgen.sh -r -a -l -n      (perform all in 1 step)
```

To configure the `mcms-noderestore` package you will need to run

`/usr/share/mcms/GenDHCPCconf.sh`. This script will generate your `/etc/dhcp/dhcpd.conf` file to serve the PXE booted NodeRestore image via tftp.

Once the certificates are generated and installed with `mcms-certgen.sh` you can proceed to configuring LDAP if desired using `/usr/share/mcms/mcms-initialize-ldap.sh`:

```
/usr/share/mcms/mcms-initialize-ldap.sh -s <nssproxy pass> -l <replication
pass> -r <ldap root admin pass> -d <base DN> -a <LDAP admin> -m <microway
pass> -c
```

Example:

```
/usr/share/mcms/mcms-initialize-ldap.sh -s Random1 -l rAndom2 -r
mostSecurePass -d dc=cl,dc=institution,dc=org -a cn=ldapadmin -m microway -c
```

Reference `/usr/share/mcms/mcms-initialize-ldap.sh --help` to see further instructions. The utility `/usr/share/mcms/make-random-password.sh` can be used to generate suitable passwords. Be sure to save these passwords for future use.

IPMImon requires configuration to monitor your cluster. See the IPMImon section for details.

Prometheus must have the head and compute nodes listed by name (not localhost) in the 'node' job section of `/etc/prometheus/prometheus.yml` and the service must be restarted afterward.

Once these steps have been completed, MCMS is ready for use. Updates are provided by the repository configured during installation.

IPMImon™ Software Configuration

This section discusses administrative issues such as configuring the automatic shutdown feature, changing passwords, and restarting the daemon software. For usage instructions, refer to the *MCMS* section in the previous section.

IPMImon Software Component Overview

There are four main software elements that interact as part of the IPMImon product:

- **IPMI Firmware** constantly runs within the BMC (baseboard management controller) on each motherboard, gathering and reporting measurement data and responding to requests.
- **IPMImon Daemon** runs on the head node, communicating with the cards, gathering data, relaying commands back to the cards, and checking for out-of-bounds measurements.
- **PHP Scripts and Apache Web Server** deliver the entire cluster's measured values over the Web and accept commands in an easy-to-use centralized interface.
- **Grafana** is a separate open-source analytics and interactive visualization application. IPMImon includes a daemon `ipmimon-exporter` that provides data to the prometheus data gathering service. Microway includes a Grafana dashboard that displays the logged IPMI data for the entire cluster over time.

Automatic Email Notification, Shutdown and Power Off

IPMImon provides the ability to automatically send an email should any measurements exceed configurable values, referred to as *alert limits*. Furthermore, IPMImon will automatically shut down and power off any nodes which exceed a second set of more stringent values, known as *critical limits*.

By default, your cluster is shipped with automatic email notification and power off disabled. You must edit the configuration files described below to configure and enable these features.

There are two files to edit, `cluster.dat` and `chassis.dat`:

- `cluster.dat` - controls global preferences of your cluster. It includes a list describing all the nodes in your cluster, including what kind of machine each is, and it specifies the email addressing preferences for automatic messages.
- `chassis.dat` - lists the sets of fans and thermal monitors to check and specifies limiting conditions that cause alerts and shutdowns for each system.

The two files are interrelated. The node list in `cluster.dat` includes a column for the chassis type referring to `$$CHASSIS` entries defined in `chassis.dat`. Both files will be located in `/etc/mcms/ipmimon`.

cluster.dat

This file lists all the nodes in your cluster. Most of the file has already been prepared for you. You need to edit the following lines:

- `$EMAIL_TO` - Provide a fully-qualified email address to send alerts to
- `$EMAIL_FROM` - Provide a fully-qualified email address from which automatic alert emails will appear to come from. This should probably correspond to a person at your site who can answer questions about the messages.
- `$EMAIL_SERVER` - Provide a single fully-qualified machine name or IP address corresponding to the local mail server at your site that will forward and deliver the messages.

Finally, you must switch on the feature by editing the following lines:

- `$ENABLE_ALERTS` – If set to no, email alerts and automatic shutdowns are disabled. Set to yes to enable them.
- `$ENABLE_SHUTDOWNS` – If set to no, email alerts can still be sent if enabled, but automatic shutdowns are disabled. Set to yes to enable them.

The subsequent lines in `cluster.dat` may be left as-is - they list the nodes in your cluster. The format is “\$NODE name chassis username password”. An example is “\$NODE node2 6028tp-htfr ADMIN ADMIN”

chassis.dat

This file governs which system sensors are checked and when an alert or shutdown condition exists. We have selected what we believe to be reasonable default values. However, the specific circumstances of your installation may require you to adjust them.

`chassis.dat` is where the specifics of the chassis are defined. While IPMI provides it's own built in thresholds, the values do not always match a customized system. Also, fans that are not working at all will not register faults as the controller sees that fan as just not being present. `chassis.dat` provides the means to set your own thresholds including expected fans. The config file itself contains many comments to guide proper configuration. Microway preconfigured this file to match your hardware. The format of each sensor is:

```
$$FAN# "Short_name" "IPMI name" "crit low" "alert low" "alert high" "crit high"
$$TEMP# "Short_name" "IPMI name" "crit low" "alert low" "alert high" "crit high"
$$VOLT# "Short_name" "IPMI name" "crit low" "alert low" "alert high" "crit high"
```

It is **very** important that `Short_name` does have any spaces in it or the web gui will not handle it properly.

To determine what IPMI sensors to monitor, you can use the `"ipmitool sensor"` command on a system that has the `"ipmi_si"` and `"ipmi_devintf"` modules loaded. That will print all sensors and their current values.

You can also find a script `"generate-chassis.sh"` in `/usr/share/mcms`. This script will generate the contents of a chassis entry from the current system by processing the `"ipmitool sensor"` output and using the system provided thresholds. You may wish to adjust the names or thresholds generated using this method.

The following is an example definition for a 6028tp-htfr Twin Squared chassis with 2 fans, 19 temperatures, and 12 voltage sensors being monitored:

```
$CHASSISTYPE 6028tp-htfr
$$FANCOUNT 2
$$FAN1 "Fan1" "FAN1" 200 500 20000 30000
$$FAN2 "Fan2" "FAN2" 200 500 20000 30000
$$TEMPCOUNT 19
$$TEMP1 "CPU1" "CPU1 Temp" 0 2 80 83
$$TEMP2 "CPU2" "CPU2 Temp" 0 2 80 83
$$TEMP3 "System" "System Temp" 0 2 60 70
$$TEMP4 "Peripheral" "Peripheral Temp" 0 2 65 75
$$TEMP5 "PCH" "PCH Temp" 0 2 70 80
$$TEMP6 "CPU1VRM" "Vcpu1VRM Temp" 0 2 85 95
$$TEMP7 "CPU2VRM" "Vcpu2VRM Temp" 0 2 85 95
$$TEMP8 "MemABVRM" "VmemABVRM Temp" 0 2 85 95
$$TEMP9 "MemCDVRM" "VmemCDVRM Temp" 0 2 85 95
$$TEMP10 "MemEFVRM" "VmemEFVRM Temp" 0 2 85 95
$$TEMP11 "MemGHVRM" "VmemGHVRM Temp" 0 2 85 95
$$TEMP12 "P1DIMMA1" "P1-DIMMA1 Temp" 0 2 65 75
$$TEMP13 "P1DB1" "P1-DIMMB1 Temp" 0 2 65 75
$$TEMP14 "P1DC1" "P1-DIMMC1 Temp" 0 2 65 75
$$TEMP15 "P1DD1" "P1-DIMMD1 Temp" 0 2 65 75
$$TEMP16 "P2DE1" "P2-DIMME1 Temp" 0 2 65 75
$$TEMP17 "P2DF1" "P2-DIMMF1 Temp" 0 2 65 75
$$TEMP18 "P2DG1" "P2-DIMMG1 Temp" 0 2 65 75
$$TEMP19 "P2DH1" "P2-DIMMH1 Temp" 0 2 65 75
$$VOLTCOUNT 12
$$VOLT1 "12V" "12V" 10.5 11.3 12.7 13.5
$$VOLT2 "5V" "5VCC" 3.5 4.3 5.7 6.5
$$VOLT3 "5VSB" "5VSB" 3.5 4.3 5.7 6.5
```

```
$$VOLT4 "3.3V" "3.3VCC" 2.5 3 3.6 4.1
$$VOLT5 "3.3VSB" "3.3VSB" 2.5 3 3.6 4.1
$$VOLT6 "Battery" "VBAT" 2 2.7 3.6 4.1
$$VOLT7 "CPU1" "Vcpu1" 1.260 1.395 1.899 2.088
$$VOLT8 "CPU2" "Vcpu2" 1.260 1.395 1.899 2.088
$$VOLT9 "DIMMAB" "VDIMMAB" 0.975 1.047 1.344 1.425
$$VOLT10 "DIMMCD" "VDIMMCD" 0.975 1.047 1.344 1.425
$$VOLT11 "DIMMEF" "VDIMMEF" 0.975 1.047 1.344 1.425
$$VOLT12 "DIMMGH" "VDIMMGH" 0.975 1.047 1.344 1.425
```

IPMImon Daemon Control

The following commands may be useful in diagnosing a suspected problem with the IPMImon hardware monitoring system. Forcing a restart of the daemon causes all nodes to be polled for new data. You need to run the commands as the root user.

To turn off the daemon that reads the data:

```
service ipmimon stop
```

To start the service:

```
service ipmimon start
```

To find out if it is running:

```
service ipmimon status
```

To restart the daemon:

```
service ipmimon restart
```

If you encounter issues with ipmimon not starting properly, stop the service with “service ipmimon stop” and then start it manually with debugging enabled with “ipmimon -d”.

Microway HPC software

Compilers

On clusters with multiple compilers installed, users will need to ensure that the include, library and execution paths are correctly set to point to the target compiler. The standard settings for popular compilers are listed below. Contact Microway technical support if you need further assistance.

Gnu Compiler Collection (GCC)

The GNU Compilers are installed and available on all Microway clusters. Microway installs gcc, g++, and gfortran on all systems. Unless otherwise requested or specified, Microway uses the packages provided by your Linux distribution for the compilers. They are installed in standard system paths and are available for all users.

Intel® oneAPI Compiler Suite

The Intel oneAPI Compiler Suite includes C, C++ and Fortran compilers in addition to multiple libraries and their MPI. By default, they are installed in `/opt/intel/oneapi` from Intel's package repository. Updates are performed along with regular system updates. The repo file can be set to disabled if you wish to update the system software but not the Intel suite. Red Hat based distributions use `/etc/yum.repos.d/oneAPI.repo` and Ubuntu uses `/etc/apt/sources.list.d/oneAPI.list`.

The Intel software may be configured for use either with their Lmod environment modules or with their `setvars.sh` script. Your cluster will be configured using one of the two methods.

When using `setvars.sh`, Microway configures entries in `/etc/profile.d` to autoload the suite using the config file `/etc/oneAPI-microway-config` as a parameter to `/opt/intel/oneapi/setvars.sh`. That config file is set to exclude unnecessary or problematic portions of the suite. Their MPI library is also excluded in the config so the choice can be managed by `mpi-selector` instead for this style configuration.

For Lmod environment modules, Intel's `/opt/intel/oneapi/modulefiles-setup.sh` is used to create `/opt/intel/oneapi/modulefiles`. That script must be rerun to update the modulefiles if the Intel oneAPI suite is updated; otherwise only the older versions will be available with `module`. That path is added to Lmod as an addition to `MODULE_PATH` enabling Lmod's `module` command to find the files it provides. When using Lmod, Microway adds the Intel compilers to the system default module set in `/mcms/modulefiles/StdEnv.lua`.

NVIDIA HPC SDK (formerly known as PGI® Compiler)

NVIDIA's HPC SDK includes both their compilers and the CUDA SDK. It is installed at `/opt/nvidia/hpc_sdk` using Nvidia's provided package repository. It can be updated along with regular system updates using `yum` or `apt`. The repo file can be set to disabled if you wish to update the system software but not the NVIDIA HPC SDK. Red Hat based distributions use `/etc/yum.repos.d/nvhpc.repo` and Ubuntu uses `/etc/apt/sources.list.d/nvhpc.list`.

Microway configures the SDK to be usable with either their provided Lmod module files found in `/opt/nvidia/hpc_sdk/modulefiles` or using an initialization script in `/etc/profile.d` when Lmod is not used.

Message Passing Interface (MPI)

MPI related files provided by Microway are found in `/usr/local/mpl`. Intel MPI is found in `/opt/intel/oneapi/mpl`. Regardless of which compiler is installed on your cluster, `mpicc`, `mpicxx`, and `mpif77` are set up to compile your code (using your chosen MPI compiler) and link it to the MPI libraries. If your compiler of choice includes a Fortran 90 compiler, `mpif90` will also be available. If your cluster was shipped with compilers in addition to the default gcc compiler, MPI has been built with each compiler.

To test MPI functionality once your cluster is operational, log in as microway and run the test script `./nas-test.sh` in microway's home directory. This test script will run continuously until `<CTRL>+<c>` is pressed. A simple cpi test is also available to be run with `mpirun`. If you have slurm, an `example-custom-job-mpl.sh` script will be present in `/home/microway` which can be submitted with `sbatch` to run a sample cpi job.

The version of MPI used in the system can be configured using either the utility `mpl-selector` or Lmod environment modules. See the MPI-Selector and Lmod sections earlier in the manual for details.

The following examples illustrate a typical procedure for building and running MPI jobs.

C Example

Log in as microway. Build the test program CPI and run it on 16 processors:

```
mpicc cpi.c -o cpi

scpf cpi /home/microway/ (this is only needed if you do not have an nfs home)

mpirun -np 16 --machinefile /etc/nodes ./cpi
```

Fortran Example

Log in as microway. Build the test program FPI and run it on 8 processors:

```
mpif77 fpi.f -o fpi

scpf fpi /home/microway/ (this is only needed if you do not have an nfs home)

mpirun -np 8 --machinefile /etc/nodes ./fpi
```

Note that the executable file has to be available to each of the nodes that will be running the job. This can be handled automatically if the directory is NFS mounted by each of the nodes. NFS mounting `/home` from the head node to all the nodes is our default setting. If you choose not to use `nfs /home`, you will need to `scpf` the files to all of the nodes.

MPI Implementations

Microway utilizes multiple different MPI libraries. Our most common configuration uses OpenMPI.

OpenMPI

OpenMPI is a highly compatible MPI library that provides a large number of accelerated communication methods. It can be used for TCP/IP and a variety of RDMA methods including InfiniBand. If your cluster includes Nvidia GPUs, Microway built OpenMPI including CUDA support.

MVAPICH2

For clusters with high speed InfiniBand networks, MVAPICH2 - a version of MPICH designed for InfiniBand - is available. Mvapich has been developed by Ohio State University as an InfiniBand specific MPI library. There are number of variants for specific purposes as well. Commercial support contracts are available for Mvapich if desired. They provide additional access to support and prebuilt optimized releases including GDR support.

MPICH4

Mpich4 is the latest version of the mpich library. Microway does not typically use mpich anymore. Please reference the upstream documentation if you have specifically requested mpich.

Intel MPI

Intel provides their own implementation of MPI with their oneAPI compiler suite. If you have requested oneAPI, Intel MPI will be available on your cluster. It will either be added to the `mpi-selector` system as an available selection or made available through Lmod environment modules.. It supports RDMA including InfiniBand and also standard TCP/IP usage.

InfiniBand and OmniPath Interconnects

Many HPC clusters include a high speed, low latency network interconnect for optimal performance. Microway provides and supports both Mellanox Infiniband and Cornelis OmniPath (formerly Intel) networks.

Users may take advantage of InfiniBand/OmniPath connectivity using applications specifically written for InfiniBand, MPI applications compiled with OpenMPI or MVAPICH, or applications using ordinary TCP/IP communications over InfiniBand (IPoIB). There is an overhead for running TCP communications over InfiniBand, so applications should be recompiled for native InfiniBand communication if at all possible.

Microway typically utilizes the open source OFED and OmniPath libraries and drivers as provided by your Linux distribution. Those software packages are updated and maintained by your standard system package manager.

In certain cases, Microway uses the binary Mellanox OFED (MOFED) packages instead to provide InfiniBand drivers, libraries, and utilities for your hardware. When using MOFED, the package repository must be updated when your OS is updated to a new point release. Please reference <https://network.nvidia.com/support/mlnx-ofed-public-repository/> to download new repository files as needed. Please ensure a new repo for your distribution's new point release prior to updating to it. There is frequently a few week delay from the time the Linux distribution's point release is made available until the MOFED release is updated to match it.

On the head node, a network management application that monitors and maintains the InfiniBand or OmniPath routing tables. Infiniband uses OpenSM which is loaded by the `opensm` or `opensmd` service. OmniPath uses the `opafm` service to provide the same functionality.

Important InfiniBand Performance Note

Network profiling with Microway's MPI Link-Checker tool has demonstrated that the OpenSM subnet manager does not automatically reconfigure the InfiniBand routing tables for optimal performance after a change has been made to the network. When the cluster is powered on, the nodes do not appear on the network at the same time, which can lead to a network imbalance. Also, any additions or subtractions of nodes from the cluster, or a reconfiguration of the InfiniBand cabling, can lead to a performance issue.

In any of these cases, the OpenSM subnet manager should be restarted with:

```
service opensm restart
```

For a cluster using OFED, you can check for IPoIB by running `ifconfig ib0`. The link speed can be checked by running `ibstat` and looking at the Rate parameter. The InfiniBand HCA should appear as a normal network adapter in all other respects. To see a list of all network devices, run `ifconfig -a`.

Slurm

Depending on the distribution used and the Slurm features required for your cluster, Slurm will either be installed from a package repository or built into packages from source. You can check the details of the installed packages with your package manager to determine the method used.

The head node runs the `slurmd` service. All compute nodes typically including the head run the `slurmd` service. These services can be controlled with `systemctl` commands. Example batch submission scripts are in the home directory for the test user microway. You can setup node and partition information in the file `/etc/slurm/slurm.conf`. In this file we have setup 4 basic partitions. The QUEUE DEFINITIONS section of `slurm.conf` is shown in the following excerpt:


```
# QUEUE DEFINITIONS

### CPU Queues ###

# Month-long jobs must be sent here (shorter jobs can be assigned higher
priorities below)
PartitionName=month-long-cpu Priority=5000 Default=NO MaxTime=31-0:00:00
State=UP Nodes=node[2] MaxNodes=6

# Week-long jobs must be sent here (shorter jobs can be assigned higher
priorities below)
PartitionName=week-long-cpu Priority=10000 Default=NO MaxTime=7-0:00:00
State=UP Nodes=node[2] MaxNodes=8

# Day-long jobs must be sent here (shorter jobs can be assigned higher
priorities below)
PartitionName=day-long-cpu Priority=20000 Default=NO MaxTime=1-0:00:00
State=UP Nodes=node[2]

# 30-Minute short, high-priority jobs may be sent here
PartitionName=short-cpu Priority=40000 Default=YES MaxTime=30:00 State=UP
Nodes=node[2]
```

You will probably want to create queues that suit your cluster needs. After making changes to the `slurm.conf` file you must distribute it to all nodes. Then run the `scontrol` command to make the changes active:

```
scp /etc/slurm/slurm.conf /etc/slurm/

scom-parallel scontrol reconfigure
```

Some helpful SLURM commands:

- To submit a batch job
 - `sbatch example_job_script.sh`
- To list job partitions and compute nodes
 - `sinfo --long`
- To list the queued and running jobs
 - `squeue --long`
- To stop (drain) a node for maintenance
 - `scontrol update nodename=example_node_5 state=DRAIN reason=Example_Reason`
- To start (undrain) a node
 - `scontrol update nodename=example_node5 state=IDLE`
- To keep a node off line for an extended period
 - `scontrol create reservation starttime=now duration=infinite user=root \
 flags=maint nodes=example_node_5`
- To view full details for a particular job
 - `scontrol show job <example_job_id_number>`
- To view job allocations within a date range
 - `sacct --allocations --starttime=2017-07-01 --endtime=2017-07-15`
- To view cluster utilization
 - `sreport cluster Utilization`

Constructing the batch submission script is fairly straight forward however there is an extensive list of possible options. You can find out more details about sbatch scripts here <https://slurm.schedmd.com/sbatch.html> and you can check the man page:

```
man sbatch
```

To get you started here is the text of a very simple example batch submission script called my_script.sh:

```
#!/bin/bash

#SBATCH --time=1

srun hostname |sort
```

In this example the time limit for the job is set to 1 minute and then the hostname command is run. Slurm by default will decide how many systems to allocate for a job based on the number of processes specified. This can be overridden with parameters in the script or specified at submission:

```
sbatch -N3 my_script.sh
```

Output of the above script submission will give you the job id:

```
salloc: Granted job allocation 12345
```

At the start of each job Slurm creates 2 files, one for output and one for errors. The files will contain the job id in the file name. To see the output of a default named job with id 12345:

```
cat slurm-12345.out

node1

node2

node3
```

An example job script with helpful comments can be found at /home/microway/example-custom-job-mpi.sh. Please reference it for additional details and helpful tips.

Microway MPI Diagnostic Suite (MMDS)

The performance of an MPI application is often limited by the speed of the slowest node in the cluster. A single slow node or intermittent link can drastically increase the run time of an MPI application! Whether you use GigE, 10GigE, 40GigE, InfiniBand or OmniPath, Link-Checker is vital for monitoring and debugging your cluster. A complimentary license for Link-Checker may already be installed on your cluster with a limited time trial license. To purchase a full license for this or another non-Microway cluster you own, please call our sales department at (508) 746-7341.

MPI Link-Checker repeatedly performs an extensive set of measurements, accurately testing latency, bandwidth, and data integrity in both directions between all pairs of nodes, and possibly even between processes within nodes. MPI Link-Checker displays its results in a detailed color-coded graphical screen.

Unique to InfiniBand networks, Microway also provides Infiniscope. It enables real time monitoring of InfiniBand hosts, switches, and connections.

MPI Link-Checker™

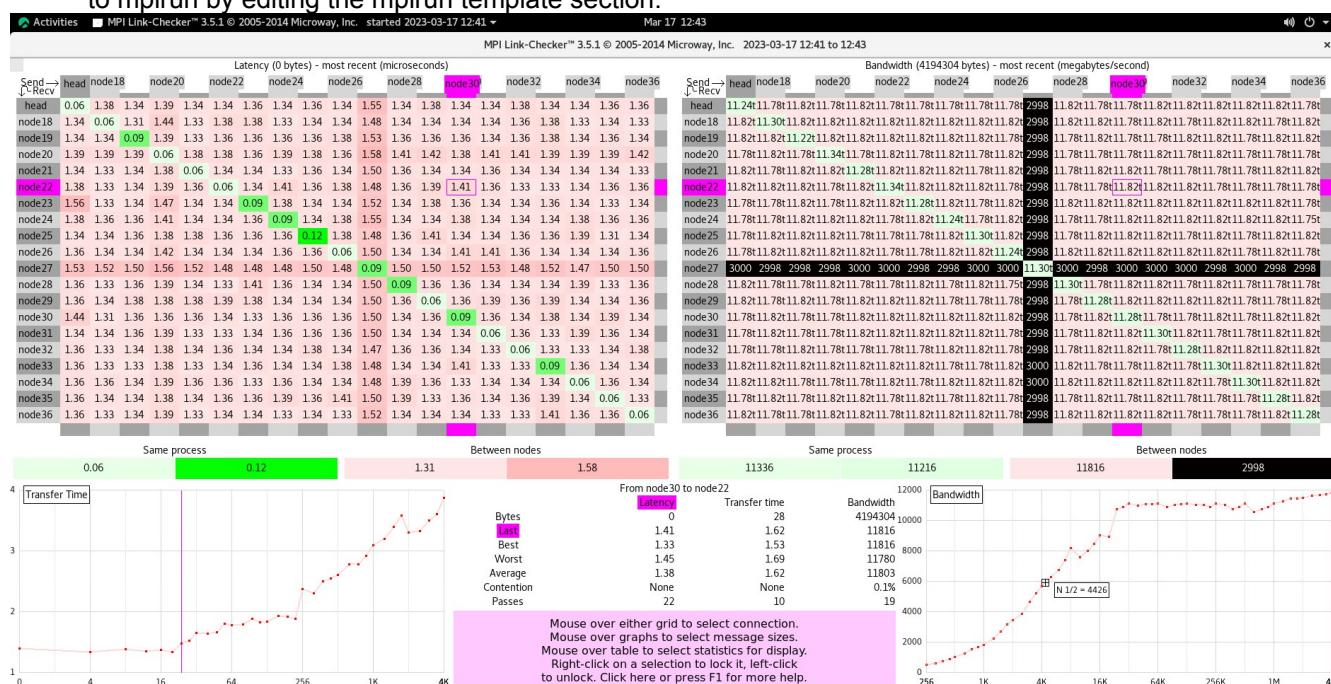
This graphical utility performs extensive bandwidth, latency and data integrity tests across the cluster network. All performance numbers and problem areas are clearly organized and labeled. Link-Checker displays a table of real-time performance numbers along with plots of latency and bandwidth against message size. Data is collected and plotted for each individual network connection.

To run MPI Link-Checker:

```
/usr/local/microway/lc -n #
```

where # is the number of cluster nodes you would like to test.

The file `/usr/local/microway/link-checker.conf` can be edited to set any required parameters to mpirun by editing the mpirun template section.



An example of the MPI Link-Checker screen is shown above. In this example, MPI Link-Checker has identified a problem with node27. You can see that the bandwidth and latency numbers reported for node8 are much worse than for other nodes. The color coding makes the problem extremely evident.

The diagonal line of green results on each table represents communication times within a single processor on each node. This measurement demonstrates the message passing speed within a single process.

The two graphs at the bottom of the screen show statistically-generated bandwidth and latency information plotted against message size for the connection that is currently highlighted with the mouse. When Link-Checker is started, these graphs will have very limited data. They will accumulate more data points and become smoother as Link-Checker continues running.

To ensure accurate measurements, do not run MPI Link-Checker while other applications are running on the cluster. Also, do not run multiple instances of MPI Link-Checker at the same time.

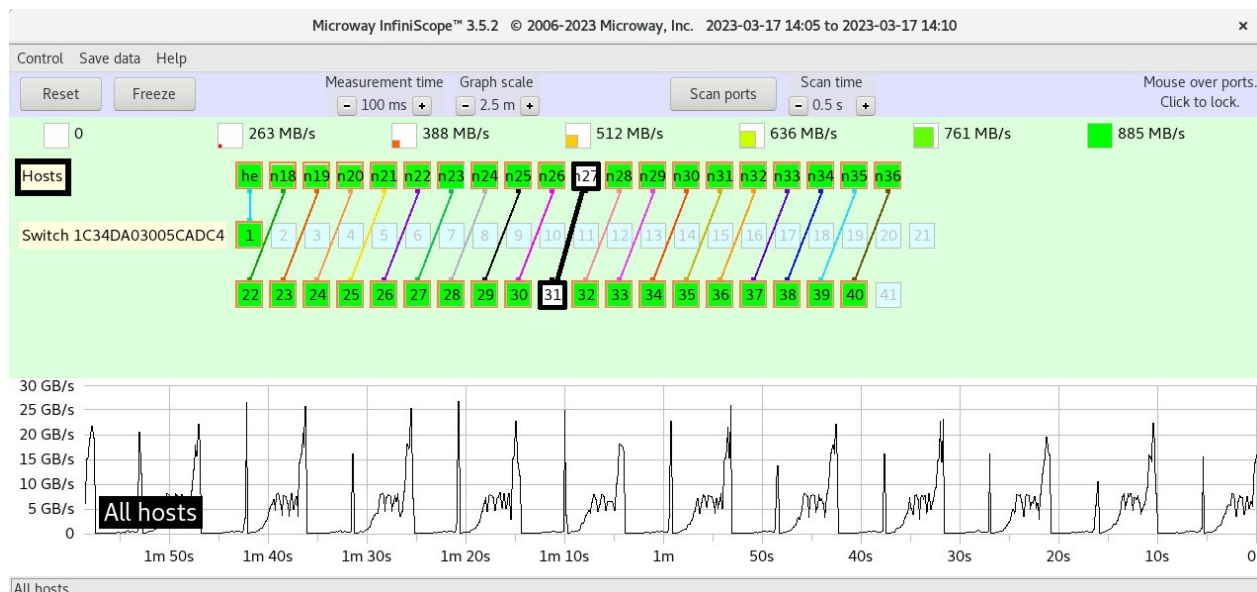
MPI performance often depends on software factors in addition to the networking hardware. It is possible to create multiple versions of the executable files using different compilers or different MPI implementations. This will demonstrate whether your latency and bandwidth issues are hardware or software related. However, you should consult the MPI Link-Checker documentation, as you will need to create additional configuration files.

Microway Infiniscope™

Microway's Infiniscope software is an InfiniBand connectivity and performance analysis tool. It presents a graphical view of all InfiniBand (IB) hosts, switches, connections, and bandwidth utilization.

To run Infiniscope:

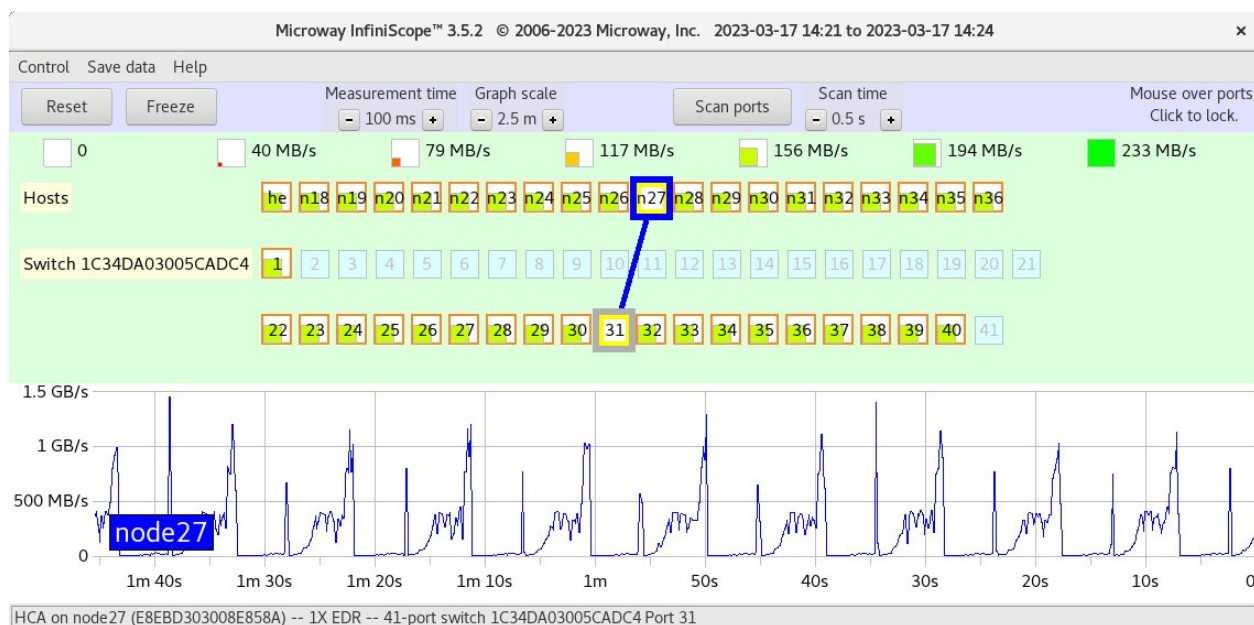
```
/usr/local/microway/is
```



18 node cluster showing all links. An issue can be seen on node27

The graphic above shows Infiniscope running on an 18 node cluster with 100Gb/s EDR InfiniBand. The color and filled portion of each node's box icon indicates the currently used bandwidth of the port. A different color link is shown between node27 and port 31 of the switch. Due to an error on that link, Infiniscope flashes the connection between black and yellow to indicate the issue.

The graph at the bottom of the screen shows IB bandwidth usage over time for the currently selected source. You can select all hosts (shown above), 1 node (shown in the following section), or for clusters with multiple IB switches, you can select all traffic per switch. Clicking "Hosts", a node box, or a switch will select that source and also provide connection information in the bottom information bar.



node27 selected showing a 1X EDR link in the information bar

In the above image, node27 has been selected to identify the issue. The information bar at the bottom of the screen shows that the link is only 1X EDR (25Gb/s). Typical IB links are 4X connections. The exception is HDR100 which uses 2 lanes (2X) at 50Gb each for a total of 100Gb/s.

If you have an issue with a reduced speed or reduced link width for an InfiniBand connection, the most likely cause is a bad cable or connection. See our troubleshooting section at the end of the manual for help diagnosing issues like this.

InfiniScope also marks any port with an InfiniBand communication error with the letter E on the port. Occasional communication errors are considered normal, but frequent errors can be a sign of an issue with the link. Please see our troubleshooting section regarding swapping cables to diagnose any communication errors on IB.

Cluster Extensions

There are several files on your Microway cluster that specify which systems (nodes) are part of the cluster. When adding additional systems into an existing cluster, these files must be updated to include the new system names. **Please follow these instructions before installing and configuring the new nodes.** Files on the head node that need to be edited are listed below.

On all clusters:

- `/etc/nodes`
 - Add new node names
- `/etc/hosts`
 - Add new IP/host entries for the regular hostname, the `-ipmi` version, and the `-ib` version (if used)
- `/etc/dhcpd.conf` or `/etc/dhcp/dhcpd.conf`
 - Change the range to accommodate the new systems. Restart the dhcpd service after this change
- `/etc/prometheus/prometheus.yml`
 - Add the new nodes to the jobname 'node' section

For clusters with IPMImon integrated into MCMS, edit `cluster.dat` in `/etc/mcms/ipmimon`.

Clusters running MPICH/MVAPICH, edit the `machines.LINUX` default hostfile in 1 of these locations:

```
/usr/local/mpich/util/machines/machines.LINUX
```

```
/usr/local/mpich/share/machines.LINUX
```

```
/usr/local/mpich/compiler/mvapich-1.2.0/share/machines.LINUX
```

Clusters with Slurm:

```
/etc/slurm/slurm.conf or /etc/slurm-llnl/slurm.conf
```

Once these files have been updated on your head node, you must distribute some of the files to the *original* nodes with `scp`. The `dhcpd`, `IPMImon`, and `MPICH` are only on the head and do not need to be distributed.

```
scp /etc/nodes /etc/
```

```
scp /etc/hosts /etc/
```

```
scp /etc/slurm/slurm.conf /etc/slurm/ (skip if not using slurm)
```

With updated configuration files present on the head node and original compute nodes, refer to the *MCMS Node Backup* section to backup an original node. This will ensure that all systems are running identical software with fully up to date configurations. When the backup is complete, restore the image onto the new (expansion) nodes. The new nodes must have the same basic type (SAS/SATA or NVME) of drive of equal or larger capacity for the image to be installed. If you wish to install the image on a smaller or different target drive, please contact Microway support for assistance with manual modifications to the image.

At the end of the restore process the new nodes will prompt for a node number. Enter the last octet of the IP address for the system to set the node number. In most cases, the last octet will match the node number. For example, node3 would typically be 10.0.0.3. If you have nonstandard name/IP pairs on your cluster, you may need to reference `/etc/hosts` on the head node when determining what number to enter at the end of the restore process.

Performance Tips and Cluster Tuning

While Microway ships clusters in a configuration ready for high computational throughput, there is no exact configuration that works best for every application. To achieve the best results with your application(s), you may need to tune your cluster. **Be aware that tuning your cluster incorrectly could greatly decrease the performance and possibly cause system/application instability.** Exercise caution when changing system values and only change one value at a time. Keep careful track of what you change to be able to undue anything in the future.

Compiler Optimization

Job execution times may be greatly decreased with advanced compiler optimization. Unfortunately, simply enabling the highest optimization flags does not always improve performance. Additionally, some applications may break or return incorrect results with certain compiler flags. Some common optimization flags are listed below - see your compiler vendor's documentation for additional options.

Compilers can generate code which is compatible with a range of processor architectures or just a single processor model (e.g., `-march=core2`, `-march=opteron-sse3`). For best performance, all applications should be compiled exclusively for your cluster's architecture. Additionally, ensure that the compiler uses all available processor extensions (e.g., `mmx`, `3dnow`, `sse`, `sse2`). See the `flags` line in `/proc/cpuinfo` for a complete list of a processor's extensions. To compile using all available optimizations from the host system performing the compilation, use `-march=native`. That will cause `gcc` to include all features present on the system. Please beware that attempting to run code built for newer hardware on older hardware may not function. When using the Intel compilers, the parameter `-xHost` performs the same functionality as `-march=native` does with GNU compilers.

For almost all compilers, the base optimization level may be set with `-O2`:

- `-O0` Disable optimization.
- `-O1` Optimize for speed and code size.
- `-O2` Optimize for speed. This is usually the default, recommended level.
- `-O3` Aggressively optimize for speed. For certain codes, may cause a decrease in performance.

Network Tuning

Our cluster benchmarks suggest that tuning the Linux kernel's network stack does not result in appreciable performance improvements on high speed network interconnects. Instead, focus on tuning your application code, selecting the proper compiler optimizations and choosing the most appropriate MPI implementation.

Extraneous Services

There are a number of system daemons and cron jobs that are automatically run by Linux. These processes will use processor and disk time that could disrupt computations, but they also provide useful services. If you are certain that you will not need these services, you can reduce the load on your cluster by disabling them.

Please do not disable a service if you don't know or don't understand what functions it provides!

For details on disabling system services/daemons, refer to the *Cluster Software* section above. To disable a cron job, simply move the cron script from the cron directory into an unused location. For example, to disable the cups cron job:

```
mkdir /etc/cron.unused
mv /etc/cron.daily/cups /etc/cron.unused/
```

Troubleshooting

This section contains a list of common cluster problems and their solutions. Some solutions require removal and installation of memory modules and CPUs. If you are not comfortable with or experienced at computer troubleshooting, or if these solutions do not help you, please contact Microway's technical support for assistance.

Linux Stress Tests

Microway has provided a set of stress tests that will ensure your system is stable. You will need to close your applications and have all other users log off the system. Log in as root and run each of these scripts:

```
/usr/local/sbin/recomp
```

```
/usr/local/sbin/rundd.sh
```

To fully stress each CPU, we provide a test script which has successfully caught many issues in the field. Running the script will initiate a process on each CPU processor core of the system. Start one instance of this CPU stress test script by running:

```
/usr/local/sbin/opteronkiller.sh
```

If the system crashes or prints any error messages while running these stress tests, look through the sections below to resolve the issues. If you see no error messages and the system continues to run the tests after a day or two of continuous tests, you can safely assume that the system is completely stable.

DOA (Dead On Arrival)

If part of your cluster is not working when you receive it, we apologize. All systems are fully configured and tested before leaving Microway, so the system was most likely damaged during shipping. Feel free to contact Microway's technical support. You may wish to ensure that the remaining portions of your cluster are working properly. The suggestions listed in this section for diagnosing problems may solve the problem you are having.

System Does Not Power On

If pressing the power button on the front doesn't have any effect, check that the switch on the power supply at the rear of the system is also on. Next, check that the power cable is firmly attached to the system and to the power strip. Ensure that the power strip is also plugged in, and that the strip's breaker has not been tripped. Finally, turn off a working system and swap power cables with the problem system. If the problem system still doesn't turn on, there is most likely a problem with the power supply or motherboard. Call technical support.

System Powers On, but Does Not Boot

If the power LED is lit and you can hear the cooling fans running, but the system is not available, directly connect a monitor and keyboard to the system. Press a few keys to wake the system up from any power saving mode. If the system wakes up and is running, refer to the *Networking Connectivity* issue below. Otherwise, power the system off and swap the memory with a system that boots properly. **To avoid damaging the systems, physically disconnect the power cable before removing or inserting memory.** If swapping memory allows the system to boot, please see the *Bad Memory* Troubleshooting section. If swapping memory has no effect, turn off the systems, unplug their power cables and swap the CPUs between the two systems. When swapping CPUs, be certain that the heatsink grease evenly covers the connection between the heatsink and the CPU in a thin layer. If swapping CPUs allows the system to boot, please see the *Bad CPU* Troubleshooting section. If the system still refuses to boot, it may need to be returned to Microway for service. Call technical support.

Network Connectivity

If the cluster loses communications with one of the nodes, it may simply be a bad cable. Verify that there is a link between the network switch and the node by checking the link lights. There should be green and/or amber lights on the switch and on the node near the network cable port. If there are no link lights, the network cable may have become loose: try re-seating the network cable on both ends. If there is still no connection, disconnect the network cable from a nearby node and swap it with the problem node's cable. If this doesn't help, use a KVM or connect a monitor and keyboard directly to the node to open a console login. You may find that the node has crashed and rebooting the node fixes the problem. If this is the case, please read the *System Crash* Troubleshooting section. If the system is still running, try moving the network cable from the current port, `eth0`, to the second port, `eth1`. If `eth1` has link lights but `eth0` does not, the motherboard may need replacement. The enclosed letter from Microway contains information on which network port should be used.

System Crash

Unfortunately, a system crash may be caused by quite a large range of issues. If there are any messages on the screen when the system crashes, be sure to write them down before rebooting the system.

One of the first items to check is the temperature of the system. If your cluster is equipped with IPMI, ensure that the CPU, motherboard and chassis temperatures are within safe margins and that all fans are spinning. If the cluster doesn't have IPMI, you may check the temperatures using the "Hardware Monitoring" section of the system BIOS. Without cool room temperatures, the cluster may not be capable of cooling itself properly - ensure your air conditioning systems are keeping the room cool. If only this particular system is overheating, remove the chassis cover and ensure that all the fans are spinning properly.

Please also refer to the *Hard Drive Errors*, *Bad Memory*, and *Bad CPU* Troubleshooting sections below.

Hard Drive Errors

Hard drive failures are an unfortunate reality when dealing with large clusters. To check for hard drive errors, look for messages similar to the following by running `dmesg` or opening `/var/log/messages`:

```
end_request: I/O error, dev sda, sector 3967232
Buffer I/O error on device sda, logical block 495904
```

All modern disks support SMART, which helps predict and detect hard drive failure. You can determine the names of the hard drive(s) in your system by running `fdisk -l`. To check the status of a disk, replace `/dev/sda` with the name of your disk and run:

```
smartctl -d ata -t short /dev/sda
```

The SMART utility will report that the hard drive will be running tests for several minutes. When the test time is over, check the status of the drive with:

```
smartctl -d ata -H /dev/sda
```

If you do not see the result `PASSED`, the drive will need to be replaced. If you have any important data on the drive, immediately copy it all to another system. Try not to reboot or power off the system, as the disk may not spin up again. You may also refer to the *Storage Administration* section above for details on hard drive monitoring.

NVME flash storage is increasingly common in HPC deployments. It offers significantly faster performance versus SAS/SATA drives, including SAS/SATA flash storage. Most `smartctl` commands work on NVME drives as well. In addition there are logs that can be queried from the drive using the `nvme cli` utility. Reference the `nvme` man page for further details.

If the system no longer boots or is unstable, but you cannot find evidence of hard drive problems, try swapping the system's hard drive with the drive of a known-working system. If the unstable node remains unstable, some other component is causing the problem. If the previously working system becomes unstable, the hard drive is faulty.

RAID Failure

When a hard drive that is part of an array fails, the RAID controller will alert you. Almost all RAID controllers will begin beeping when an array is compromised by a drive error or drive failure. You may also see kernel messages reporting the details of the problem. **Please take care when dealing with this situation: contact technical support before taking any steps with which you are unfamiliar or uncertain.**

You first need to identify what state the array is in and what caused the problem. **Do not yet reboot the system!** If the RAID management tools are still operating, use them to check the status. Otherwise, directly connect a monitor, KVM, or serial console to the system to determine the RAID status. Our technical support staff may not be able to resolve the issue until all the details of the failure have been discovered.

You will need to count the number of drives in the array and note which drive numbers are no longer members of the array. You should also note the status of the "hot spare" drive and determine if the controller has started rebuilding the array using the "hot spare" drive.

In most cases, a single drive has failed and the RAID controller will automatically rebuild the array using the “hot spare” hard drive, if available. In this case, the RAID management tools will report that the array is rebuilding and will display the rebuild progress. Use the management tools to verify which drive is faulty and physically remove it. The drive will need to be returned to Microway for a replacement. When the new drive is inserted, it should be configured as the new “hot spare” hard drive.

If the system crashed, rebooted, or the array is no longer accessible from Linux, the problem may be more complex. Please keep in mind that the drives are not being written to while the motherboard BIOS and RAID controller BIOS are initializing. Once the operating system has booted and the RAID controller drivers are loaded, the drives will be available for writing. In almost all cases, you will be able to resolve the situation from the RAID BIOS, so you should keep the system from booting into the OS until you are certain the array is in a safe condition. This will ensure that each drive in the array will be kept “in sync” with the others in the event that more advanced recovery methods are necessary.

If you are remote to the system via an ssh session there are command line tools that can help you assess the situation.

- For LSI
 - Use the command `storcli64`
 - If you do not have the `storcli64` command loaded on your system, go to <http://www.broadcom.com> and search `storcli64` to find the download.
 - some useful status commands:
 - `storcli64 /c0 show all`
 - This command will show all current status information about your controller and RAIDs for controller 0. Some systems, especially larger storage systems, may have multiple controllers. If you are unsure of how many controllers you have run the command
 - `storcli64 show all`
 - This will show you a list of controllers.
 - `storcli64 /c0 show pdfailevents`
 - This command zeros right in on if you have any physical drive failure events logged
 - `storcli64 /c0/vall show all`
 - This command shows you all virtual drives
 - `storcli64 /c0/v0 show all`
 - This command shows you information specific to the first virtual drive
 - if you add the `-h` switch to the of a command `storcli64` will show you the available options. For instance:
 - `storcli64 /c0 show -h`
 - This will show you the options and variations available for the show command on controller 0.
- For Areca
 - use the command `cli64`
 - If you do not have the `cli64` command loaded on your system, go to <http://www.areca.us/support/> to find the application.
 - Running “`cli64 help`” will display the available options for the utility
 - Some useful status commands:
 - To see the hardware and firmware version for your Areca RAID controller
 - `cli64 sys info`
 - To see information about your virtual drives
 - `cli64 vsf info`
 - To see individual disk information
 - `cli64 disk info`
 - To see event logs
 - `cli64 event info`

You may also refer to the *Storage Administration* section above for details on RAID usage and monitoring.

Bad Memory

EDAC reporting

Current generation Intel and AMD systems support the kernel feature EDAC (Enhanced Detection and Correction). This provides detailed information in the event of a memory error. This information is only presented in the `dmesg` and the `/var/log/messages` file. It should provide a MC#, channel #, and row#. An example error on the 2nd CPU is:

```
kernel: EDAC MC3: 0 CE memory read error on CPU_SrcID#1_MC#1_Chan#1_DIMM#0
(channel:1 slot:0 page:0x30b89f2 offset:0xd40 grain:32 syndrome:0x0 -
err_code:0x0000:0x009f socket:1 imc:1 rank:0 bg:1 ba:3 row:0xa05 col:0x70)
```

Knowledge of the memory controller is required to decode the MC and Chan portions into a DIMM location. The output from the command “`edac-util -v`” can be helpful. Indices start from 0. In the above example, the 2nd memory controller(MC#1) in the 2nd CPU(CPU_SrcID#1) had an error on it's 2nd memory channel(Chan#1) on the 1st DIMM(DIMM#0) of the channel. For that system the error occurred on memory slot CPU2 DIMMD1. Please note that some motherboards label CPUs and DIMMs starting from 0 and others start from 1. If trying to work on the 1st CPU memory slots, to ensure the right location, check if anything is labeled CPU0 before working on CPU1. To positively identify your error, please contact Microway with the EDAC message and “`edac-util -v`” output and we can identify exactly which DIMM is reporting the error.

MCELOG reporting

Older systems utilized the utility mcelog to report memory errors. If you suspect malfunctioning memory, there are several courses of action. The Linux kernel will log any non-fatal memory errors as Machine Check Exceptions (MCE) and will print the message `Machine check events logged` to `dmesg` and into the file `/var/log/messages`. The mcelog daemon will then decode the error and log it into `/var/log/mcelog`. Here is an example entry from that file:

```
MCE 0
CPU 0 BANK 7
MISC 140080a86 ADDR 1ee11d4a40
TIME 1409957568 Fri Sep 5 18:52:48 2014
MCG status:
MCi status:
Corrected error
MCi_MISC register valid
MCi_ADDR register valid
MCA: MEMORY CONTROLLER RD_CHANNEL0_ERR
Transaction: Memory read error
STATUS 8c00004000010090 MCGSTATUS 0
MCGCAP 1000c1b APICID 0 SOCKETID 0
CPUTID Vendor Intel Family 6 Model 62
```

CPU and memory channel indices start at 0 in the reports. The information reported can be used to narrow down the problem to that particular CPU and channel of RAM. The above example shows the error on the first memory channel of the first CPU. With 1 DIMM per channel the bad DIMM is obvious (1st DIMM on 1st CPU), but if more than 1 DIMM is installed per channel you'll need to swap memory from that channel with DIMMs from other channels to isolate the specific failing DIMM once another error occurs.

Corrected vs Uncorrected errors

EDAC “CE” (corrected error) and mcelog “Corrected” memory errors should only be considered a problem when multiple errors are occurring on a given system. The developers of mcelog typically consider 10 corrected errors per day acceptable. Any report of “UE” or “Uncorrected” memory errors should be considered failed hardware which requires repair.

Isolating the problem by swapping components

To ensure the memory is the real problem, now that you know which physical CPU and potentially which DIMM is having trouble, power off the system, unplug it and remove the chassis cover. **To avoid damaging the system, physically disconnect the power cable before removing or inserting memory.** Locate the proper CPU and remove the 4 or 8 memory modules adjacent to the CPU. Swap them with the memory modules on one of the other CPUs and then boot the system. If you know which DIMM is reporting the error, note which slot it is swapped with. Once the system has booted, if using mcelog, log in as root and clear the MCE log by running:

```
mv /var/log/mcelog /var/log/mcelog.old

service mcelog restart
```

Systems using EDAC reporting do not require any log or service changes.

Then run the *Linux Stress Tests* described at the beginning of this *Troubleshooting* section. Unless there was a seating or other similar issue that swapping the memory resolved, the kernel will eventually report additional errors, as the bad components will still be causing errors. If the memory was causing the errors, they will now be reported on a different CPU (in the location the expected bad DIMM was swapped to, if known), because you physically moved them. You should contact Microway technical support to obtain replacements for these memory modules. If the errors are still being reported on the same CPU, the memory was not causing the errors - refer to the *Bad CPU* section below.

If the memory errors are fatal, the system will crash before any error logging is possible. In these cases, it is simplest to swap all the memory between the problematic system and a functional system and then run the *Linux Stress Tests* described at the beginning of this *Troubleshooting* section. If the problematic system continues to crash, see the *Bad CPU* section below. If the problematic system becomes stable and the functional system begins crashing, contact Microway technical support to obtain replacement memory modules.

Bad CPU

If you have not already tested the memory, please refer to the section *Bad Memory* above. Memory and CPU problems often display identical symptoms, but memory failures are much more common.

A faulty CPU may exhibit itself by causing Machine Check Exceptions (MCE) or Enhanced Detection and Correction (EDAC) errors as described above in the *Bad Memory* section, causing system crashes or even preventing the system from properly booting. The most straightforward troubleshooting method is to swap all the CPUs between the problematic system and functional systems and then run the *Linux Stress Tests* described at the beginning of this *Troubleshooting* section. Swap each CPU into a different system if possible to isolate which component is bad quickly. **To avoid damaging the systems, physically disconnect the power cable before removing or inserting CPUs.** When swapping CPUs, if exposed in the process, be certain that the heatsink grease evenly covers the connection between the heatsink and the CPU. If the problematic system continues to display unwanted behavior, the motherboard may need replacement - contact Microway technical support. If the problematic system becomes stable and the functional system or systems begin displaying unwanted behavior, contact Microway technical support to obtain replacement CPUs.

MPI Problems

Problems Starting MPI Jobs

- If you receive the error “command not found” when running `mpirun`, you need to check the `PATH` environment variable. The MPI `bin/` directory, which is usually in `/usr/local/`, should be listed in the `PATH` (e.g., `/usr/local/mpich/bin/`, `/usr/local mpi/gcc/openmpi-4.1.4/bin/`).
- The error “Child process exited while making connection to remote process” usually indicates that some of the compute nodes were not able to access the MPI executable file that was passed to `mpirun`. Ensure that all nodes have a copy of the executable file(s).
- Mismatched compile time and runtime MPI libraries can cause a variety of strange errors to occur. Ensure you are running your MPI job with the same MPI library it was built with.
- Messages about being unable to lock memory when using InfiniBand indicate your account does not have a high enough `memlock` ulimit. Running “`ulimit -l`” will show the current value. The limit is set by `/etc/security/limits.conf` but can also be limited by Slurm or a systemd configuration. The `gnome-terminal` user service imposes a low `memlock` ulimit. Microway provides a fix for this on our clusters as shipped.
- Messages similar to “error while loading shared libraries” or “cannot open shared object file” suggest that the nodes do not have access to, or cannot locate, a library needed by the application. Ensure that there is a copy of the library on the nodes and that the environment variable `LD_LIBRARY_PATH` contains this library directory. You can also build the library into the executable at compile time by using the compiler flag `-static`, but this will increase the size of the executable.

MPI Jobs Exit Prematurely

- The most common cause for MPI jobs to unexpectedly die is a bug in the MPI application. Try running the same application with a few different sets of input data. If the application only hangs for one particular set of data, the problem is almost certainly the application itself. Your best course of action will be contacting the software developer, as debugging parallel applications is complex. If the application seems to be unstable for all inputs, there may be a problem with the MPI library or network drivers.
- A crashed or rebooted node, or a disconnected cable, could also cause your application to die. Ensure that all nodes are available and that the network communications are still running smoothly.

Basic troubleshooting rules to live by:

- Record and read through any error messages.
 - If your computer is stuck on a boot screen, or is in some way unresponsive to keyboard commands, then write down the full text of any error messages. If you are in a facility that allows cell phones then snap a picture of the output. This can be very helpful information as you proceed through the process of solving the issue.
- Review the basic error logs for your OS.
 - For Linux
 - Run the command `dmesg` and read through the output.
 - More often than not, the information you need in order to understand what is happening is logged there.
 - Check syslog messages
 - `journalctl` is used on most modern distributions
 - `/var/log/messages` may also have additional history in some cases
 - on Ubuntu `/var/log/kern.log` contains kernel log history
 - For Windows
 - Review the System and Application Logs in the Windows Event Viewer.
 - Where to find event logs is different for different Windows versions.
 - Make sure you take some time (while things are working and everyone is calm) to find your way to the event logs and write down or remember how to get there.
 - If you're not sure, use the search bar and type event.
- Write down or record the following:
 - What changed recently on your system.
 - If you have updated your system or installed a new application or device driver it may have impacted the functionality. When things break and before doing anything else think about what's changed and write that down. Seemingly unrelated changes can sometimes be the cause of issues.
 - What jobs/tasks was the computer running before the breakdown.
 - Are there any audible alarms or led indicators showing error.
 - What steps have you taken to troubleshoot the problem.
 - ** This is important whether you call us or not. Sometimes in the heat of battle it's easy to lose track of the steps you have completed or solutions you've tried. This can quickly add up to you losing your way. Better to go slow and record then go fast and get lost.
- Google, Bing or DuckDuckGo (pick your search engine) are your friends.
 - The fastest way to find manuals that are up to date for your system is to search on-line.
 - Many times others have encountered the same error you are experiencing. Their solution may be helpful to you. Use good judgment here though because believe it or not, sometimes people write things that are not true on the internet!

We hope that your cluster provides you with years of good service in pursuit of knowledge. We value our role in helping you in this endeavor. Call us at 508-746-7341 or send email to tech@microway.com if you need help.