Chapter 2

CM-5 and Scale 3 System Startup Procedures

This chapter describes the procedure for bringing a CM-5 or a Scale 3 from a powered down condition to the state where it is ready to partition (described in Chapter 3). The only procedure specific to machine type is the one for powering up; this is because of cabinetry differences.

The procedures described in this chapter are:

- powering up the system
  - CM-5 (other than Scale 3) (see Section 2.1)
  - Scale 3 (see Section 2.2)
- checking the environmental sensor board's status panel (see Section 2.3)
- updating system files (see Section 2.4)
- preparing for SDA bringup and CM-5/Scale 3 partitioning (see Section 2.5)
- initializing the CMIO bus system (see Section 2.6)
- bringing up the SDA (see Section 2.7)

These procedures are presented in several levels of detail, from a high-level view of the general tasks to detailed descriptions of each step.

- Figure 3 presents in flowchart format the major tasks involved in bringing up a CM-5 or Scale 3 system.
- Figure 4 presents in quick-reference format the individual steps involved in each bringup task.

Copyright © 1995 Thinking Machines Corporation

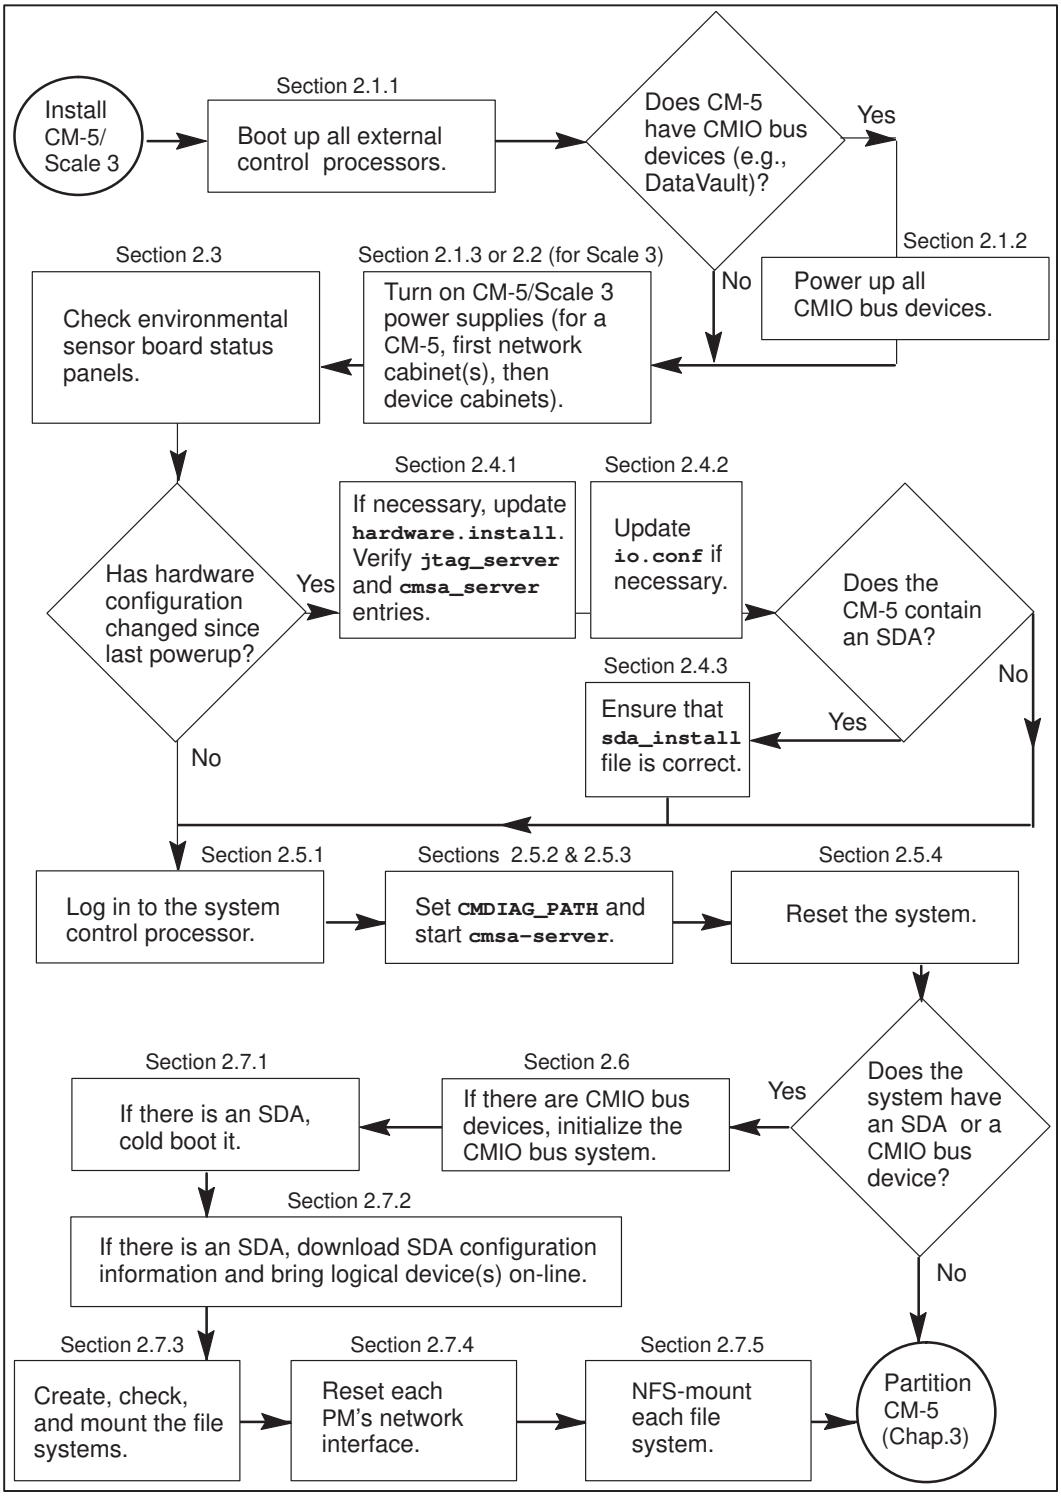

Figure 3. Major task summary of the basic CM-5 startup procedure.

Copyright © 1995 Thinking Machines Corporation

1. Boot up each external control processor, if there are any.
2. If the system includes CMIO bus devices (e.g., DataVault), power up all I/O devices.
3. Power up the Scale 3 cabinet / CM-5 cabinets (first the network cabinet[s], then the device cabinet[s]). Check the status panel of each cabinet's environmental sensor board.
4. Update the **`/etc/cm/configuration`** system files, if necessary: **`hardware.install`**, **`io.conf`**, and **`sda.install`** (check logical devices).
5. Ensure that **`jtag_server`** and **`cmsa_server`** are set in **`hardware.install`**, or that the corresponding environment variables are set.
6. Log in to the system control processor as root.
7. Set **`CMDIAG_PATH`** to specify the diagnostic library pathname (the default is **`/usr/diag/cmdiag`**).
8. Start **`/usr/etc/cmsa-server`** on a CP that has an SNI board. Use the **`-r`** switch when appropriate.
9. Run **`/usr/etc/cmreset`** on the system control processor and **`cmreset -s`** on each of the other CPs and IOCPs.
10. If the system includes a CMIO bus device(s), execute **`/usr/etc/io_cold_boot`**.
11. If the system includes an SDA:
    (1) Execute **`sda_config`** on each IOCP.
    (2) Execute **`sda_cold_boot`** on one IOPC, and **`sda_cold_boot -f`** on any other IOCPs.
    (3) Execute the **`sdacc`** command **`sc`** once for the system.
    (4) On each IOCP, execute the **`sdacc`** command **`lld`**, and execute the **`sdacc`** commands **`sld`** and **`online`** once for each of the IOCP's logical devices.
    (5) On each IOCP, for each logical device, execute
    `sdafsck /dev/rsda`*logcl_dev_num*`a`
    `mkdir -p` */mount_point*
    `mount -t sdafs /dev/rsda`*logcl_dev_num*`a` */mount_point*
    (6) On each IOCP, execute **`exportfs -a`**.
    (7) On each CP that is not an IOCP, execute **`cmreset -s`**.
    (8) On each CP (including IOCPs), NFS-mount each file system: execute **`mount sdafs`** *IOCP_name*`:/`*mount_point* `/`*local_mount_point*, where *IOCP_name*`:/`*mount_point* is the mount point for the file system on the IOCP managing it, and `/`*local_mount_point* is the mount point on the CP that is NFS-mounting the file system.

Figure 4. Scale 3 / CM-5 startup procedure.

Copyright © 1995 Thinking Machines Corporation

2.1 Powering Up a CM-5 System

This section describes how to power up a CM-5 system (other than a Scale 3). Powering up a Scale 3 system is described in Section 2.2.

The power-up procedure assumes that the CM-5 is completely installed (hardware and software), including all cabling.

2.1.1 Boot External Control Processors

Power up any external control processors and verify that their boot sequence is successful. The location of the power switch depends on which Sun model is used to implement the control processor. If you have any questions about this step, refer to the applicable Sun documentation.

2.1.2 Power Up All Peripherals

If the CM-5 system includes peripheral devices that reside on a CMIO bus, such as a DataVault, CM-HIPPI, CM-IOPG, or VMEIO devices connected to a VMEIO host computer, apply power to their power supplies and boot up their station managers. If you need detailed instructions, please refer to Part III of this manual.

2.1.3 Power Up the CM-5, Including Its SDA

Note that since the SDA resides inside the CM-5 cabinetry, it is automatically powered up when the CM-5 is powered up.

As described in Chapter 1, the CM-5 system comprises one or more device cabinets and one or more network cabinets (one network cabinet for each four device cabinets); each cabinet in the system is equipped with its own set of power switches. Power up the network cabinet(s) first, to allow their master clock(s) time to stabilize before the device cabinets come on-line (the device cabinets' clocks reference the network cabinets' master clock). The powerup procedures are described below.

When the processing nodes power up, their boot-mode sequence is indicated on the LED panel by a left-to-right "chase pattern." After applying

Copyright © 1995 Thinking Machines Corporation

power to the CM-5 cabinets, check each cabinet's LED panel to verify that it has powered up successfully. If this pattern is not displayed, check to see that the cabinet's LED switch is in the "fast copy" mode. Table 3 summarizes the various LED mode settings for this switch (the location of the LED switch is shown in Figure 6).

Table 3. CM-5 LED mode switch settings.

Hex Value	Mode	Definition
0*	Freeze LEDs	Maintains current state of LEDs.
1	Fast Copy (default)	Copies PN values to LEDs; special boot-mode sequence generates left-to-right chase pattern.
3	Interleaved Copy	Displays PN values on LEDs; groups of four even-numbered rows shift their display to the left, and groups of four odd-numbered rows shift to the right.
5	Random	Generates random patterns.
7	Interleaved Random	Same shifting behavior as mode 3 but with values supplied by random number generator.
9	LEDs On	Turns all LEDs on.
A	LEDs Off	Turns all LEDs off.
B	Blink	Alternately turns all LEDs on and off; on for 0.5 second and off for 0.5 second.
C	Test	Continually runs power-up self-tests; these tests are described in the *CM-5 Field Service Guide*.
D	Display PM Loop	Used in hardware diagnostic sessions to trace connectivity problems.

* Switch settings 0, 2, 4, 6, 8, E, and F all specify Freeze mode.

Copyright © 1995 Thinking Machines Corporation

Powering Up a Network Cabinet

Figure 5 shows the location of a network cabinet's power switches and the layout of the circuit breakers. To reach these switches, slide the covering panel to the right.

Turn on each network cabinet's power in the following sequence:

- main circuit breaker
- PDU (power distribution unit) controller
- contactor
- LEDs
- +5 V supplies (any order)
- +2 V supplies (any order)

Powering Up a Device Cabinet

On device cabinets, the power switches are located behind the louvered corner panel that covers the cabinet's power supplies. See Figure 6. The panel is held closed by magnetic latches along the main face of the cabinet and is hinged on the cabinet's end wall. To open the panel, briefly press against the latched face and then release; the panel swings out away from the cabinet, exposing the power supply bay. Again, see Figure 6.

Apply power to the device cabinet power supplies in the following sequence. Figure 7 shows the locations of the referenced circuit breakers. Repeat this sequence for all device cabinets in the system.

- main circuit breaker — CB1
- AC and DR/CN — CB3 and CB4
- +5 V supplies — CB5, CB7, CB9, CB11, CB13, CB15, CB17, and CB19
- +2 V supplies — CB6, CB8, CB10, CB12, CB14, CB16, CB18, and CB20
- LEDs — CB21 and CB22

Copyright © 1995 Thinking Machines Corporation

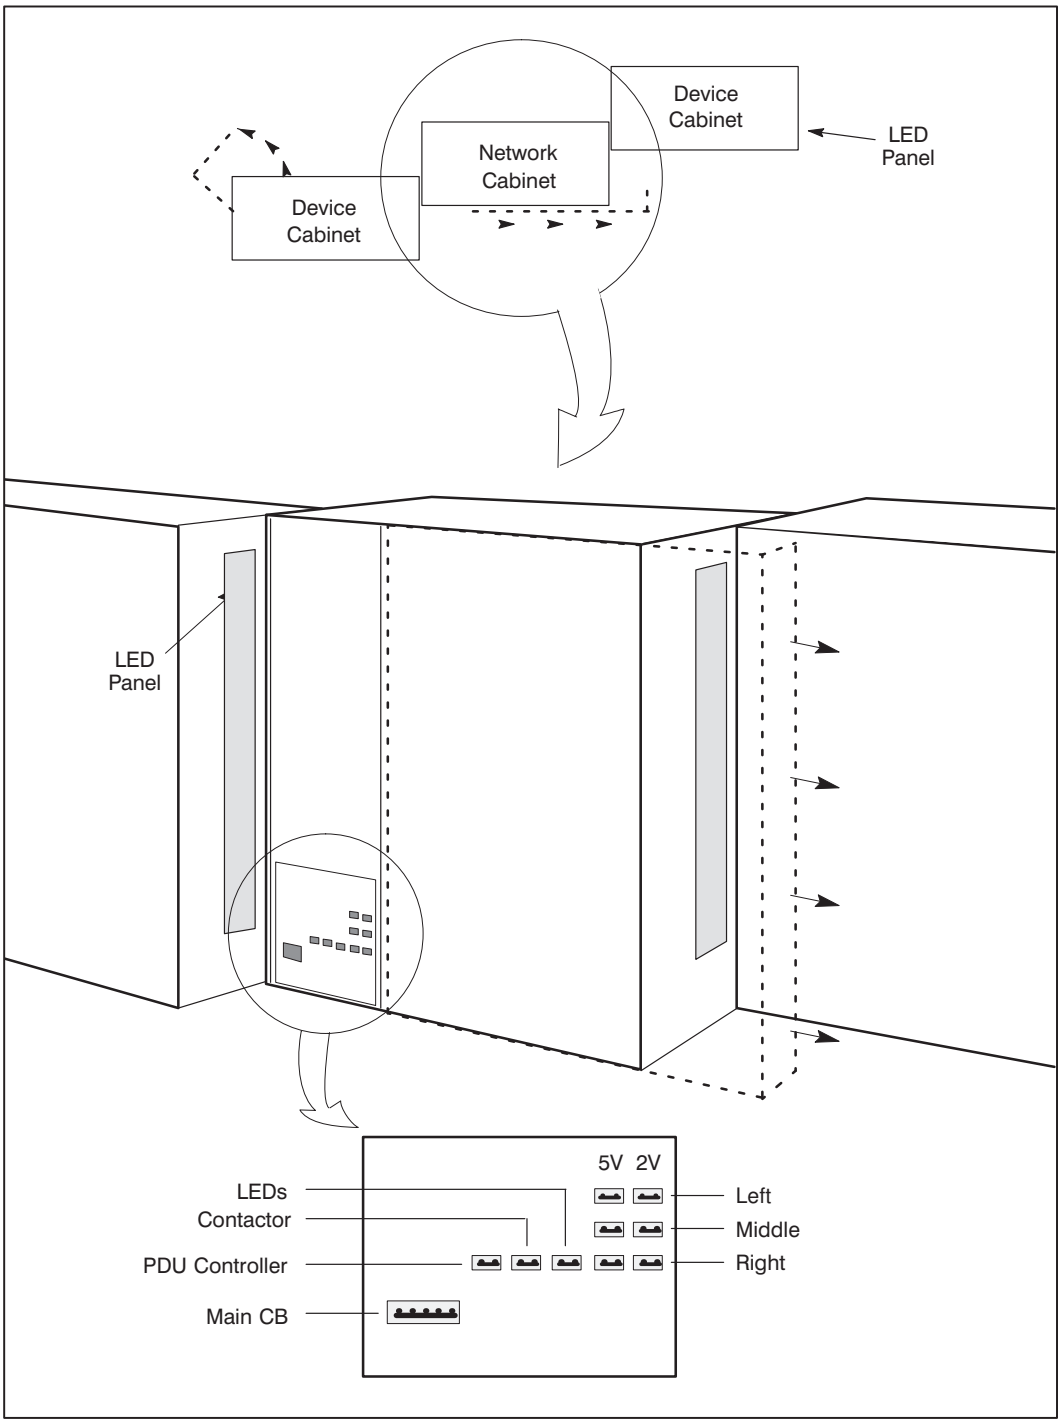

Figure 5. Access to network cabinet circuit breakers.

Copyright © 1995 Thinking Machines Corporation

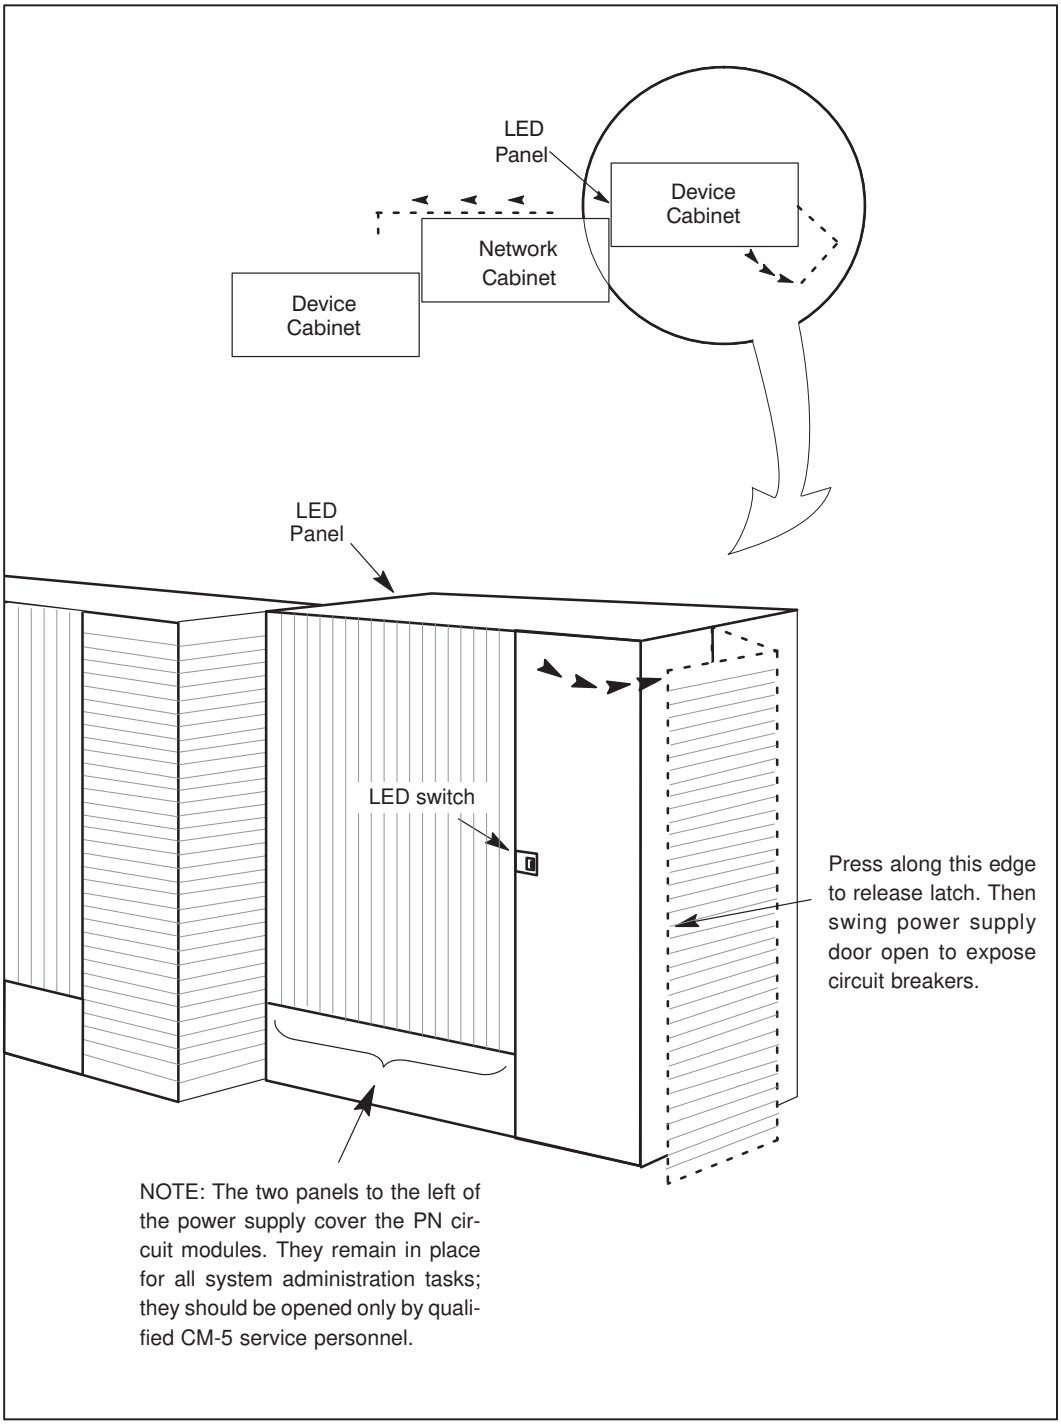

Figure 6. Access to device cabinet circuit breakers.

Copyright © 1995 Thinking Machines Corporation

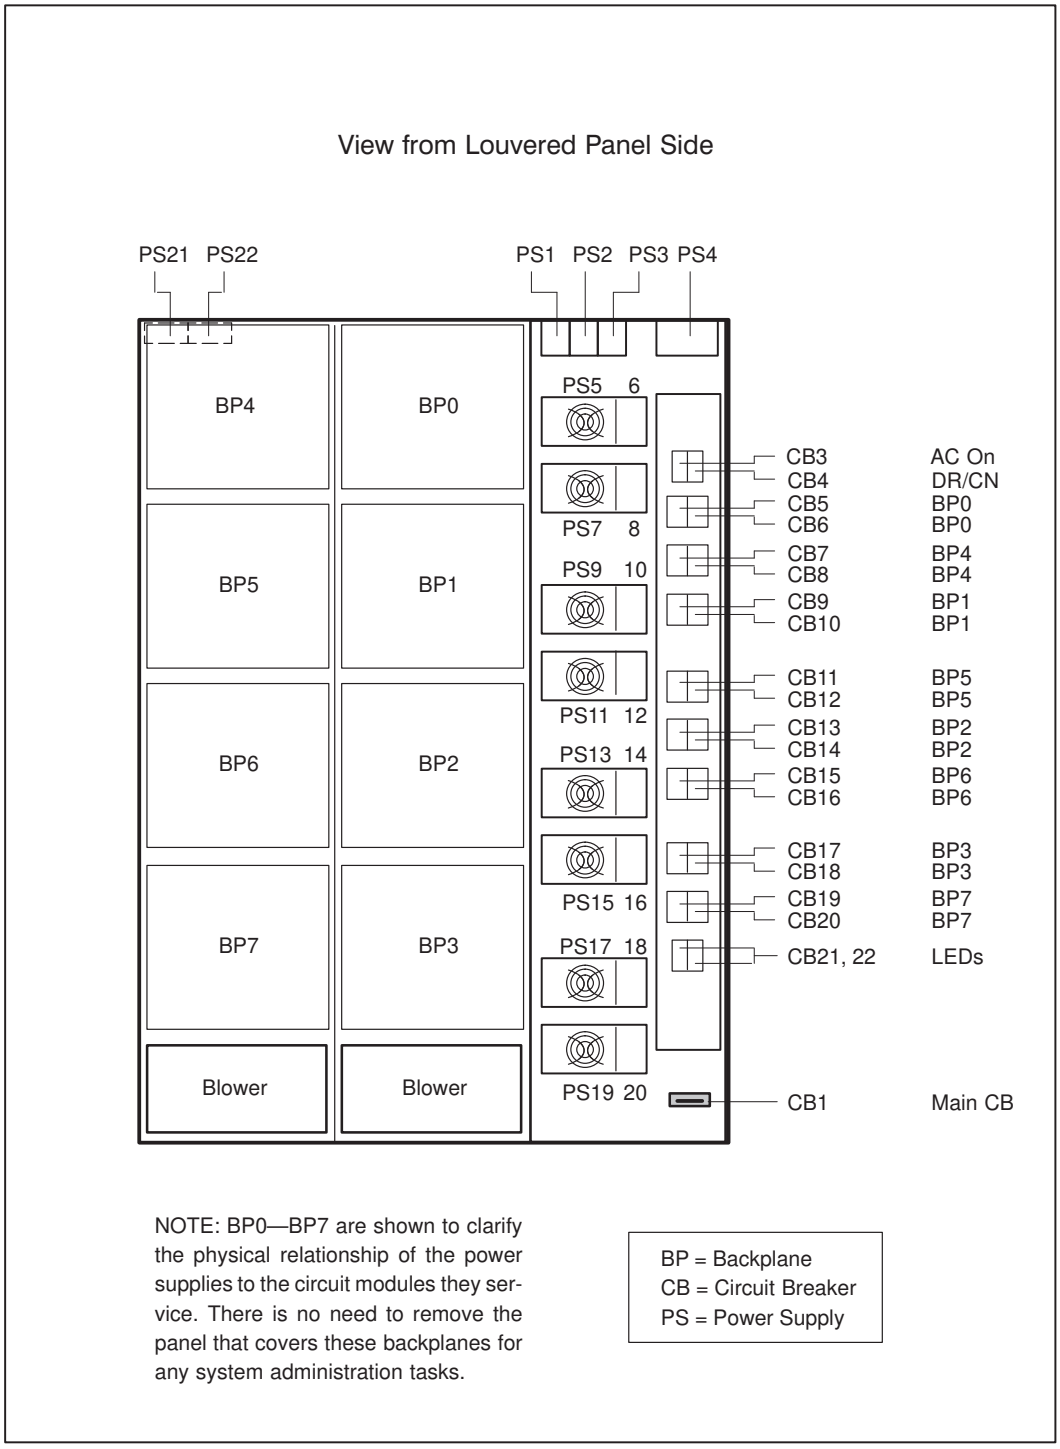

Figure 7. Device cabinet circuit breaker locations.

Copyright © 1995 Thinking Machines Corporation

2.2 Powering Up a Scale 3 System, Including Its SDA

This section describes how to power up a Scale 3 system. (Powering up a CM-5 [other than Scale 3] system is described in Section 2.1.) Note that since the SDA resides inside the Scale 3 cabinetry, it is automatically powered up when the Scale 3 is powered up.

The powerup procedure assumes that the Scale 3 is completely installed (hardware and software), including all cabling.

The Scale 3's main circuit breaker is located behind the louvered panels at the back of the machine. To open the panels, press the release lever at the top of the panels, causing them to slide apart a few inches. Slide the panels further apart by pushing them outward as far as they will go. See Figure 8.

When the panels have slid apart, you can see the circuit breakers near the bottom of the machine, at the lower left corner of the power distribution unit, as shown in Figure 8 and detailed in Figure 9. Switch on the main circuit breaker by pushing up the breaker handle.

Now verify that the smaller breaker switches to the right of the main switch are all on (in the up position) — they may or may not have been shut off during the last shutdown procedure. If any of the smaller breaker switches are not on, turn them on by pushing up the breaker handles. The order in which these are turned on is arbitrary.

After applying power to the CM-5 Scale 3, check the LED window on the front of the machine to be sure the LEDs are displaying their factory-programmed pattern — a left-to-right "chase pattern." This pattern indicates that the processing nodes have powered up.

NOTE: The LED panel is purposely disabled in CM-5 Scale 3 systems installed at European sites, in compliance with the European safety code standard.

Copyright © 1995 Thinking Machines Corporation

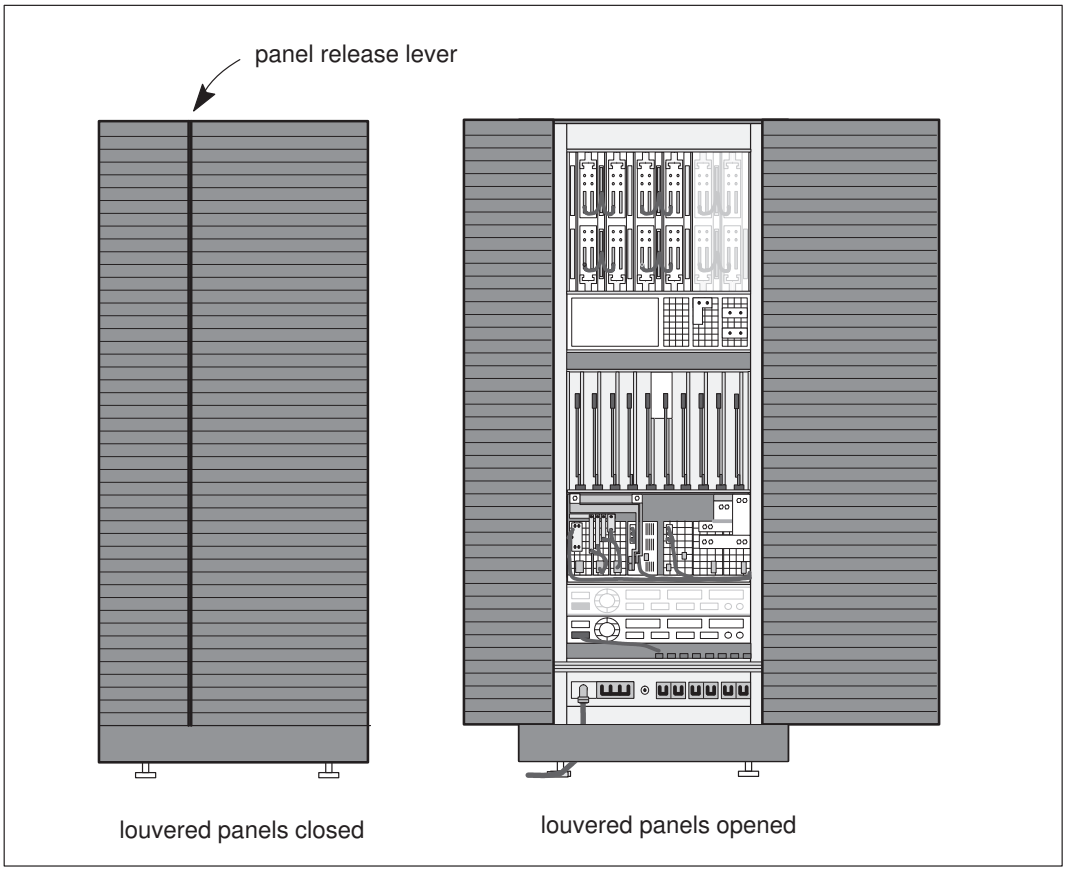

Figure 8. CM-5 Scale 3, back view. To open the back panels, press the panel release lever and manually push apart the two panels.

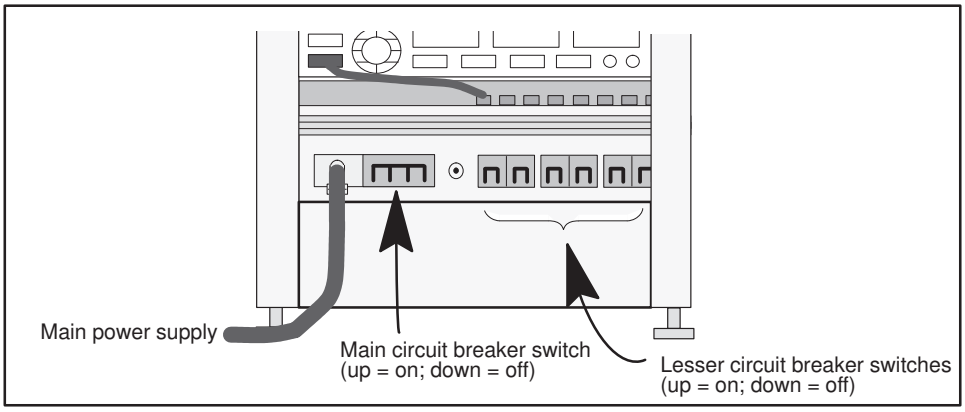

Figure 9. CM-5 Scale 3, back view (partial, with louvered panels removed), showing detail of circuit breaker switches.

Copyright © 1995 Thinking Machines Corporation

2.3 Check the Environmental Sensor Board's Status Panel

NOTE: As of this writing, CM-5 and Scale 3 systems in the field are being upgraded with this new environmental sensor board; your system may not yet contain this board.

In a Scale 3, the sensor board's status panel is located in the back of the cabinet, between the SDA backplane and the PN backplane; in a CM-5 cabinet, the sensor board is in the same bay as the power supplies, near the upper PDU.

When you have finished powering up the CM-5 or Scale 3, check each cabinet's environment sensor board's status panel to be sure its power supplies are operating correctly and that the air flow and temperature in and around the machine is within specification:

- All the magnetic indicators should be off — black, not bright green.
- All the green LEDs should be lit; all the red LEDs should be dark. (Only one LED set is visible on the status panel of a CM-5 (other than Scale 3) cabinet — the rest are intentionally hidden by sheet metal because they are not functional on a CM-5.)

If you find one or more magnetic indicators on, i.e., green, or one or more red LEDs lit upon powerup, these may be leftover signs of a remedied problem. The machine may have been shut down automatically because one of the sensors tripped. If this is the case, reset the status panel by pressing the switch at the far right of the panel, wait three minutes, and recheck the status panel; all the magnetic indicators and red LEDs should remain off.

For more information about the status panel and what its indicators mean, see Section 6.3 in Chapter 6, "Preventive Maintenance."

Copyright © 1995 Thinking Machines Corporation

2.4 Update the System Files

The system administrator is responsible for keeping the three main system files, listed below, up to date. There is only one copy of each file per CM-5 system; it resides on the system control processor.

- **/etc/cm/configuration/hardware.install** (see Section 2.4.1) must be present on every CM-5 system. It is created at the factory to define the system's hardware configuration. As long as the system's hardware configuration remains unchanged, this file requires no attention.
- **/etc/cm/configuration/io.conf** (see Section 2.4.2) must be present on every CM-5 system other than Scale 3 that includes CMIO bus devices, such as a DataVault. **io.conf** is created at the factory to define the configuration of the CMIO bus(es) and all devices connected to it. As long as the system's CMIO bus configuration remains unchanged, this file requires no attention.
- **/etc/cm/configuration/sda.install** (see Section 2.4.3) must be present on every CM-5 system that includes an SDA. **sda.install** is created at the factory to define the SDA's disk configuration. This file may be modified by the administrator to change the logical device configuration or to make logical disk substitutions.

2.4.1 Ensure that hardware.install Is Correct

All CM-5 systems must contain an accurate **/etc/cm/configuration/hardware.install** file. The **cmsa-server** program verifies the **hardware.install** entries against the hardware it knows to be present in the CM-5. If **hardware.install** is not up to date, therefore, **cmsa-server** will not start (and a message will be logged in **cmsa-server.log**).

If any CM-5 hardware components have been installed, removed, or repositioned since the last time the system was powered up, update **/etc/cm/configuration/hardware.install** to reflect the changes. One-to-one hardware replacements do not require editing the file since the net change is zero.

Appendix A describes the **hardware.install** contents for CM-5 systems (other than Scale 3); Appendix B describes the **hardware.install** contents for Scale 3 systems. If you are still uncertain about how to edit this file,

Copyright © 1995 Thinking Machines Corporation

please contact your Thinking Machines Corporation representative for guidance.

The jtag_server and cmsa_server entries

We call attention to the following **hardware.install** entries:

- **jtag_server =** *"CP_name"*

  The **jtagserver** process runs on the system control processor, the CP that contains the Diagnostic Network hardware (SDN board). The **jtag_server** entry must contain the hostname of this CP.

- **cmsa_server =** *"CP_name"*

  The **cmsa-server** process runs on a CP that contains the Control Network and Data Network (data router) hardware (SNI board), which may or may not be the system control processor. The **cmsa_server** entry must contain the hostname of the CP on which the **cmsa-server** is running.

The **jtag_server** and **cmsa_server** entries in **hardware.install** eliminate the need to set the environment variables **CMSA_SERVER** and **JTAG_SERVER**. If, for some reason, the **cmsa_server** and **jtag_server** entries are not made in **hardware.install**, the corresponding environment variables *must* be set. If the **cmsa_server** and **jtag_server** entries are present in **hardware.install** and the environment variables are also set, the setting of the environment variables overrides the setting of the **hardware.install** entries.

Copyright © 1995 Thinking Machines Corporation

Entries Specific to the CM-5E

If your site's CM-5 is being upgraded to CM-5E status, the machine's **hardware.install** file needs to be updated.

New PN_Backplane Entries. For any PN backplane that has been upgraded with CM-5E hardware, find the corresponding **PN_Backplane** section in **hardware.install**, and make the following changes:

- If a backplane contains all CM-5E components (SuperSPARC PNs, the new CM-5E NIs, and the new CM-5E Data Network (DR), change the backplane's **PN** entry to **SSPNFGT**. Note that the **PN** entry might be in a definition at the beginning of the file, currently; if so, you can either edit the old definition, or — if the old definition is still needed, as in the case of a partial machine upgrade — you can create a new definition.

  If you change the **PN** entry to **SSPNFGT**, remove that backplane's **PNMEM** entry (for backplanes without vector units) or **PNDP** entry (for backplanes with vector units)

- If the boards in the backplane have the new Data Network:
  - Change any **SPI** entry to **SPIGT**.
  - Change any **FILLER** entry to **FILLERGT**.

New DR_Backplane Entries. If a DR backplane has been upgraded with CM-5E hardware, find the corresponding **DR_Backplane** section in **hardware.install**, and make the following changes:

- If the boards in the backplane have the new DR:
  - Change any **DR** entry to **DRGT**.
  - Change any **DR5** entry to **DR5GT**.

Copyright © 1995 Thinking Machines Corporation

2.4.2 For CM-5 (Other than Scale 3) Sites Only: Update io.conf if CMIO Bus Device Configuration Has Changed

This section applies only if your system contains one or more CMIO bus devices, such as a DataVault. If your system contains no CMIO bus devices, **/etc/cm/configuration/io.conf** need not be present.

io.conf defines the bus attributes, such as station ID and bus arbiter status, of the I/O devices connected to the CMIO bus.

If the system's CMIO bus device (e.g., DataVault) configuration has changed since the system was last powered up, edit **/etc/cm/configuration/io.conf** to incorporate those changes. If the change also involves adding, removing, or relocating any IOBA (input/output bus adapter) hardware, you need to edit **hardware.install** as well.

Appendix C describes **io.conf**. If you are uncertain about how to edit this file, please contact your Thinking Machines Customer Support representative for guidance.

2.4.3 Systems with an SDA

This section applies only if your system contains an SDA. If your system does not contain an SDA, **/etc/cm/configuration/sda.install** is not present.

sda.install is created at the factory by executing the **sda_install** command, which automatically generates the biggest module of the file, **DISKS**. Two smaller parts of the file — **SPARES** and **LOGICAL_DEVICES** — may or may not have been incorporated at the factory. It is important at this time to examine your system's **sda.install** to ensure that it contains a **LOGICAL_DEVICE** module, and that it is correct, because this module is necessary for the SDA bringup procedure. Section 11.3 in Chapter 11 describes how to create or edit this module.

Copyright © 1995 Thinking Machines Corporation

2.5 Prepare for CM-5 Partitioning and for SDA Bringup

This section describes the steps necessary to prepare the CM-5 for partitioning and SDA bringup. Even if your system does not include an SDA, you must perform all the tasks in this section.

To describe how to perform the tasks, we will use a sample system to illustrate the procedures in this section and the next. Figure 10 shows our sample system, which comprises the following:

- A system control processor, **syscon**. Our sample system's system control processor is a CP that contains both an SNI board and an SDN board. (Not all system control processors contain both boards. Although the system control processor must always contain the SDN board, in some systems a different CP contains the SNI.) In many CM-5 systems, the system control processor serves also as a partition manager, but for simplicity, this one does not.
- Two CPs used as partition managers, **beethoven** and **haydn**.
- Two IOCPs used to manage file systems, **iocp1** and **iocp2**.
- An SDA configured with three logical devices. The file system for **logical device 1** will be managed by **iocp1**; the file system for **logical device 2** will also be managed by **iocp1**; the file system for **logical device 3** will be managed by **iocp2**.

2.5.1 Log In to the System Control Processor

Log in to the system control processor as root:

```
syscon% su
passwd: password
syscon#
```


Copyright © 1995 Thinking Machines Corporation

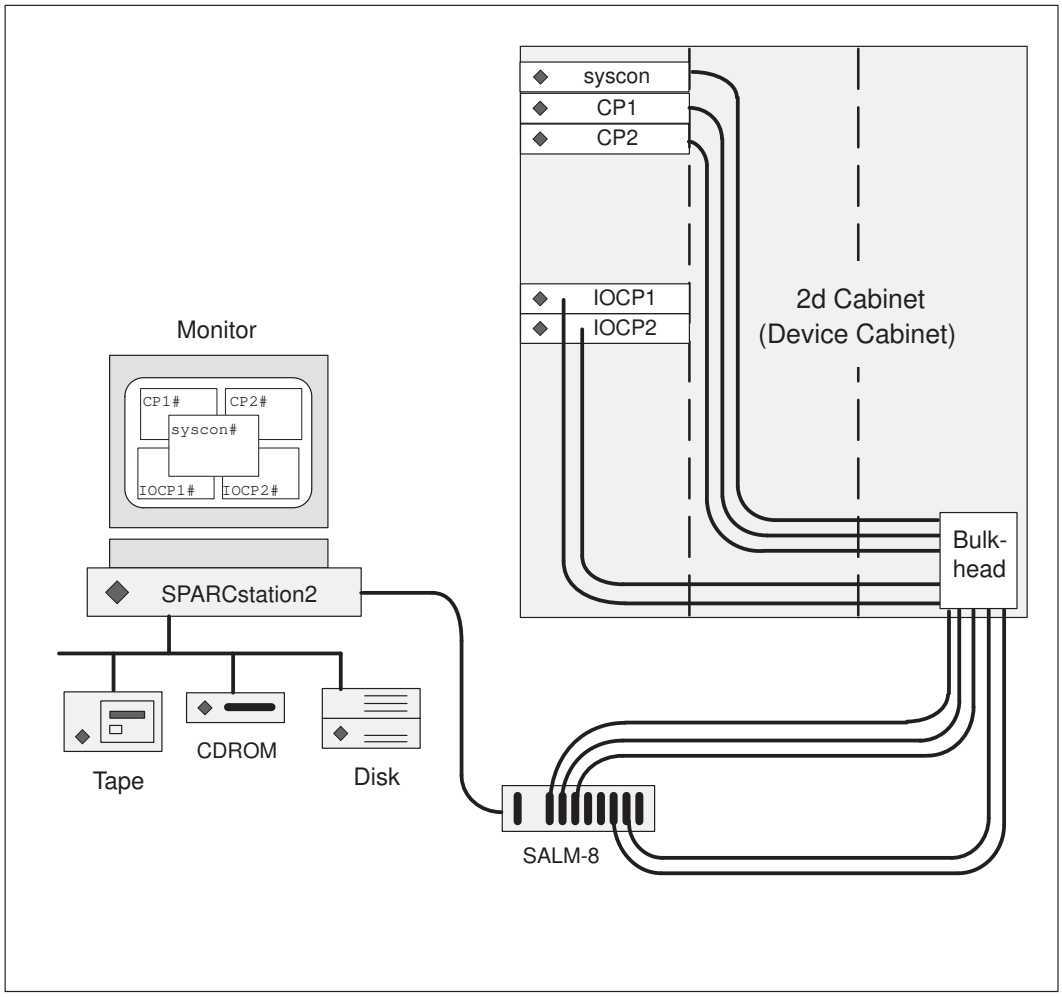

Figure 10. A representation of the sample system being brought up (showing CPs and IOCPs).

2.5.2 Set the CMDIAG_PATH Environment Variable

Set the **`CMDIAG_PATH`** *pathname* environment variable, where *pathname* is the name of the directory containing the diagnostic utilities (usually **`/usr/diag/cmdiag`**). For example,

```
syscon# setenv CMDIAG_PATH /usr/diag/cmdiag
```


Copyright © 1995 Thinking Machines Corporation

2.5.3 Start cmsa-server

Start the **cmsa-server** program running on a CP that contains the Control Network and Data Network hardware (SNI board). This CP may be the system control processor, as it is in our sample system:

```
syscon# /usr/etc/cmsa-server -r
```

cmsa-server must be running whenever the CM-5 is partitioned or reset. The **-r** switch is used *only* when starting **cmsa-server** after powerup to ensure that any configured partitions are set to the proper initial state, i.e., **ALLOCATED**. (If the timesharing daemon had previously exited abnormally, the partition may have been erroneously left in an **ACTIVE** state.) Do not use the **-r** switch when any partitions are truly active.

2.5.4 Reset the CM-5 and Individual CP Network Interfaces

The **cmreset** command is required after each powerup cycle to synchronize system clocks and initialize all registers and switches to a known state. (For a CM-5E system, CMSA automatically sets the clock source to 40 MHz. The clock speed of a CM-5 system is automatically set to 32 MHz.) To reset the CM-5 hardware, perform the following steps:

1. Run **/usr/etc/cmreset** on the system control processor to execute a system-wide hardware reset.

   cmreset by default uses 256 words as the system-wide default FIFO size for the CM−5E Data Networks at Height 3 and above. Note that with a 256-word FIFO size and a small timesharing quantum, the time to do context switching can become prohibitive. If you want to specify a different FIFO size for the system, include **-f** *FIFO_size* on the **cmreset** command line; *FIFO_size* can be 16, 64, 128, or 256.

2. On each other CP (including any IOCPs), execute **cmreset -s** to reset its network interface (you don't need to execute **cmreset -s** on the system control processor because **cmreset** now does an implicit **cmreset -s**).

   cmreset -s resets the network interface of the CP on which it is executed.

Copyright © 1995 Thinking Machines Corporation

On our sample system, which has two control processors and two IOCPs in addition to the system control processor, system reset is as follows:

```
syscon# /usr/etc/cmreset
syscon# rsh beethoven /usr/etc/cmreset -s
syscon# rsh haydn /usr/etc/cmreset -s
syscon# rsh iocp1 /usr/etc/cmreset -s
syscon# rsh iocp2 /usr/etc/cmreset -s
```

2.6 For CM-5 (Other than Scale 3) Sites Only: Initializing the CMIO Bus System

This step is necessary only if the CM-5 system includes one or more CMIO bus devices, such as a DataVault.

After resetting the CM-5, download the I/O kernel to a processor in each IOBA (the CMIO bus interface) by executing:

```
syscon# /usr/etc/io_cold_boot
```

io_cold_boot is executed only once per CM-5 system. It must be executed from the system control processor.

Usually, no switches are required. However, if the files that **io_cold_boot** uses are not in the default locations, you must specify one or more switches:

- The **-I** *filename* switch specifies the location of **io.conf**, which usually resides in **/etc/cm/configuration**.
- The **-K** *filename* switch specifies the location of **io_kernel**, which usually resides in **/usr/etc**.
- The **-B** *filename* switch specifies the directory that contains the **cmboot** and secondary boot code (default is **/usr/etc**).

For more detailed information about **io_cold_boot**, see its manual page in Appendix J.

Copyright © 1995 Thinking Machines Corporation

2.7 Bringing Up an SDA

When the steps in the previous sections have been completed, you can bring up the system's SDA. If the system has no SDA, you can begin partitioning the machine, as described in the next chapter.

Before beginning the tasks in this section, be certain that the SDA's logical device(s) are correct as they appear in the LOGICAL_DEVICES module of **sda.install**, as mentioned in Section 2.4.3.

The following bringup procedure assumes that the steps in Sections 2.1, 2.4, 2.5, and 2.6 — and only the steps in these sections — have been completed.

Recall that, as described in the beginning of Section 2.5, our sample system contains one SDA configured with three logical devices, two partition managers, and two IOCPs. We will bring up three file systems on our system: two will be managed by **iocp1**, and one will be managed by **iocp2**.

2.7.1 Cold Boot the SDA

To cold boot the SDA:

1. Update the VMUNIX kernel on the SDA's IOCP(s) with the current SDA configuration, based on **/etc/cm/configuration/sda.install**, by executing **sda_config** (with no switches) on each IOCP that will manage one or more file systems for the SDA. For example, on our sample system:

   ```
   syscon# rsh iocp1 /usr/etc/sda_config
   syscon# rsh iocp2 /usr/etc/sda_config
   ```

   For more detailed information about **sda_config**, see its manual page.

2. Download and boot the SDA's IOP kernel by executing **sda_cold_boot** on each of the SDA's IOCPs. Execute **sda_cold_boot** with no switches on one IOCP. If the SDA has any other associated IOCPs, execute **sda_cold_boot** with the **-f** switch. For example, on our sample system:

   ```
   syscon# rsh iocp1 /usr/etc/sda_cold_boot
   syscon# rsh iocp2 /usr/etc/sda_cold_boot -f
   ```


Copyright © 1995 Thinking Machines Corporation

The **-f** switch marks all the disk storage nodes as booted, without actually booting them.

Usually, no other switches are required. However, if the files that **sda_cold_boot** uses are not in the default locations, you must also specify one or more switches:

- The **-B** *pathname* switch specifies the directory that contains the **cmboot** and secondary boot code (default is **/usr/etc**).
- The **-K** *filename* switch specifies the kernel executable file (default is **/usr/etc/scn_kernel**).

For more information about **sda_cold_boot**, see its manual page in Appendix J.

The cold-boot process takes several minutes. For example, a 120-disk SDA takes approximately six minutes to cold boot.

2.7.2 Download the SDA Configuration Information and Bring the Logical Device(s) On-Line

Figure 11 shows the procedure for downloading the SDA configuration information and bringing the SDA's logical device(s) on line. To illustrate the general steps listed in Figure 11, we present the command sequence below, a continuation of the bringup procedure for this chapter's sample system.

Before beginning the actual **sdacc** command sequence, again recall that our sample system's **sda.install** file lists three logical devices. We have decided that the file system for logical device **1** will be managed by **iocp1**, the file system for logical device **2** will also be managed by **iocp1**, and the file system for logical device **3** will be managed by **iocp2**.

Once this preliminary information has been decided, we can begin the **sdacc** command sequence, as shown below:

```
syscon# rsh iocp1 /usr/etc/sdacc
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: sc
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: lld
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: sld 1
<LD 1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: online
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: sld 2
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: online
```

(cont'd)

Copyright © 1995 Thinking Machines Corporation

```
<LD 2, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: quit
syscon# rsh iocp2 /usr/etc/sdacc
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: lld
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: sld 3
<LD 3, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: online
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: quit
syscon#
```

1. For each logical device to be brought on-line, determine its *logical device number* (an integer corresponding to the logical device's relative position in **sda.install**) and decide which IOCP will manage its file system.

2. Invoke the **sdacc** program:

   ```
   syscon# rsh IOCP_name /usr/etc/sdacc
   <LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>:
   ```

3. DO THIS STEP ONLY ONCE FOR THE *ENTIRE SYSTEM*; that is, do this step only one time on only one IOCP, regardless of how many IOCPs and logical devices are in the system.

   Download the SDA configuration information by executing the **sdacc** command **sc** (**set-config**), as shown below:

   ```
   <LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: sc
   <LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>:
   ```

   The **sc** command requires the file **/etc/cm/configuration/sda.parameters** to be present. If **sc** fails with a message that this file is not present, contact your Thinking Machines Corporation Customer Support representative.

4. DO THIS STEP ONLY ONCE FOR *EACH IOCP*; that is, if an IOCP will manage more than one logical device, do this step for only the first logical device it will manage.

   Load the logical device by executing **lld** (**load-logical-device**), which takes no arguments, as shown below:

   ```
   <LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: lld
   <LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>:
   ```

(cont'd)

Figure 11. The procedure for downloading SDA configuration information and bringing logical devices on-line (page 1 of 2).

Copyright © 1995 Thinking Machines Corporation

5. Bring the IOCP's logical device(s) on-line. For each logical device associated with the IOCP, perform the following two steps:

   a. Execute **sld** (**select-logical-device**), which takes one argument, *logical_device_number*. *logical_device_number* is the number relative to the logical device's position in **sda.install** (further explained in Section 11.3 of Chapter 11).

   b. Execute the **online** command, which takes no arguments.

   For example:

```
<LD -1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: sld 1
<LD 1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: online
```

6. Based on your system's current state, take one of the following three courses of action:

   (1) If there are no more logical devices to bring on-line:

   Quit out of the **sdacc** program:

```
<LD 1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: quit
```

   Continue with Section 2.7.3 ("Check and Mount Each File System on Its IOCP").

   (2) If there are more logical devices to bring on-line, but none of their file systems will be managed by the IOCP on which **sdacc** is currently running:

   Quit out of the **sdacc** program:

```
<LD 1, CAB -1, BP -1, SLOT -1, SCN NI -1, DISK ID -1 -1>: quit
```

   Continue with step 2, above, on the next IOCP.

   (3) If the IOCP on which **sdacc** is currently running will manage another logical device that is yet to be brought on-line:

   Do not quit out of the **sdacc** program. Continue with step 5, above, for the next logical device to be managed by the current IOCP.

Figure 11, cont'd. The procedure for downloading SDA configuration information and bringing logical devices on-line.

Copyright © 1995 Thinking Machines Corporation

2.7.3 Check and Mount Each File System on Its IOCP

NOTE

The procedure described in this section assumes that the SDA file systems have already been created. If the file systems have not yet been created, create them now; to create a file system, execute the following command on the IOCP that will manage it:

syscon# **rsh** *IOCP_name* **sdafsmake /dev/rsda***N***a** \
[*number_of_data_drives*]

rsda*N***a** specifies the logical device to be mounted, where *N* is the number of the logical device (further explained in Chapter 11, Section 11.3).

For *number_of_data_drives*, an optional argument, specify the number of data drives **rsda***N***a** contains (excluding parity drives); this information is in **/etc/cm/configuration/sda.install** (again, see Chapter 11). If you do not specify this argument, **sdafsmake** queries the operating system to determine this information.

Finally, verify that **rsda***N***a** is present in **/dev**.

Now that each logical device has been brought on-line, you can check its file system and mount it for use by executing three commands and editing a file, as listed below. The commands must be executed once for each logical device/file system.

(1) **/usr/etc/sdafsck** *logical_device_name*
Checks SDA file system consistency.

(2) **mkdir -p /***mount_point*
Makes the **/***mount_point* directory and creates any missing parent directories.

Copyright © 1995 Thinking Machines Corporation

(3) **/usr/etc/mount -t sdafs /dev/sda***N***a /***mount_point*
Attaches the file system **/dev/sda***N***a** to the file system hierarchy at the place specified by the **/***mount_point* argument. (The **-t sdafs** argument specifies the type of the file system.)

(4) If the file **/etc/exports** already exists on that IOCP, add **/***mount_point* to it and save the file.

If the file **/etc/exports** does not yet exist on that IOCP, create it, put **/***mount_point* in it, save the file, and execute **rpc.mountd -n** and **/etc/nfsd 8 &**.

After executing the three commands indicated above for each logical device/file system, a fourth command, **/usr/etc/exportfs -a**, must be executed once on each IOCP that hosts one or more logical devices/file systems. The command **/usr/etc/exportfs -a** makes all the pathnames listed in **/etc/exports** available to NFS clients for mounting over the network.

For example, to check and mount a file system for each logical device brought on line in Section 2.7.2, execute the following commands:

```
syscon# rsh iocp1 sdafsck /dev/rsda1a
syscon# rsh iocp1 mkdir -p /mount_point
syscon# rsh iocp1 mount -t sdafs /dev/sda1a /mount_point
syscon# rsh iocp1 sdafsck /dev/rsda2a
syscon# rsh iocp1 mkdir -p /mount_point
syscon# rsh iocp1 mount -t sdafs /dev/sda2a /mount_point
<Create and/or edit /etc/exports on iocp1.>
syscon# rsh iocp1 exportfs -a
syscon# rsh iocp2 sdafsck /dev/rsda3a
syscon# rsh iocp2 mkdir -p /mount_point
syscon# rsh iocp2 mount -t sdafs /dev/sda3a /mount_point
<Create and/or edit /etc/exports on iocp2.>
syscon# rsh iocp2 exportfs -a
```

Optional: Using /etc/fstab

You can set up each IOCP's **/etc/fstab** to include an entry for its SDA file system.

For example, **iocp1**'s **/etc/fstab** SDA file system entry would look like this:

```
/dev/sda1a   /mount_point     sdafs  noauto,rw 0     0
```


Copyright © 1995 Thinking Machines Corporation

The SDA file system would be mounted on **iocp1** by executing:

```
syscon# rsh iocp1 sdafsck /dev/rsda1a
syscon# rsh iocp1 mkdir -p /mount_point
syscon# rsh iocp1 mount /mount_point
syscon# rsh iocp1 exportfs -a
```

2.7.4 Reset Each PM's Network Interface

IMPORTANT

Do not execute **cmreset -s** on a CP functioning as an IOCP; that is, do not execute **cmreset -s** on a CP that hosts an SDA file system. Execute **cmreset -s** only on a CP functioning as:

- the system control processor only
- a partition manager (PM) only
- both the system control processor and a PM

Reset the system control processor's NI and each partition manager's NI by executing **cmreset -s**. Be careful not to execute **cmreset** with no switches! On our sample system, the reset would be performed as follows:

```
syscon# /usr/etc/cmreset -s
syscon# rsh beethoven /usr/etc/cmreset -s
syscon# rsh haydn /usr/etc/cmreset -s
```


Copyright © 1995 Thinking Machines Corporation

2.7.5 NFS-Mount Each File System

NFS-mount each SDA file system on each CP. For example,

```
syscon# mount -o intr -t sdafs iocp1:/mount_point  \
/local_mount_point
syscon# mount -o intr -t sdafs iocp1:/mount_point  \
/local_mount_point
syscon# mount -o intr -t sdafs iocp2:/mount_point  \
/local_mount_point
syscon# rsh beethoven mount -o intr -t sdafs \
iocp1:/mount_point  /local_mount_point
syscon# rsh beethoven mount -o intr -t sdafs \
iocp1:/mount_point  /local_mount_point
syscon# rsh beethoven mount -o intr -t sdafs \
iocp2:/mount_point  /local_mount_point
syscon# rsh haydn mount -o intr -t sdafs \
iocp1:/mount_point  /local_mount_point
syscon# rsh haydn mount -o intr -t sdafs \
iocp1:/mount_point  /local_mount_point
syscon# rsh haydn mount -o intr -t sdafs \
iocp2:/mount_point  /local_mount_point
```

IMPORTANT

Note that these file systems are NFS-mounted interruptible; that is, the **mount -o intr** flag *is* included on the command lines. (In the previous release of CMOST, Version 7.3, we recommended that NFS-mounted file systems be mounted *un*interruptible because of a software restriction that has now been removed.)

Optional: Using /etc/fstab

If you want to set up each CP's **/etc/fstab** to NFS-mount the SDA file systems, you must include the **noauto** option in each **/etc/fstab** SDA file

Copyright © 1995 Thinking Machines Corporation

system entry. The **noauto** option must be included because the SDA is not yet up when the CPs boot.

For example, a CP's **/etc/fstab** SDA file system entry would look like this:

IOCP_name:/*mount_pt* /*local_mount_pt* **nfs noauto,rw,intr 0 0**

The SDA file system would be mounted on the CP by executing

syscon# **rsh** *CP_name* **mount -o intr /***local_mount_point*

Copyright © 1995 Thinking Machines Corporation