LogicLoader™ User's Manual

(LogicLoader Version 2.4)

Logic Product Development
Published: April 2008

REVISION HISTORY

REV	EDITOR	REVISION DESCRIPTION	LoLo Ver.	APPROVAL	DATE
A	Szilveszter Balogh, Jed Anderson	Started with manual for LoLo 2.3 (PN 70000016) with the following changes: - Added Section 1.6: LogicLoader Labs - Section 5.2: updated 'burn' command - New Section 5.3: added 'dd' command - Section 5.4: updated 'erase' command - Removed Section on booting from NAND - New Section 10 to describe partitions - New Section 11 to describe File Systems - Section 12: updated for new approach to YAFFS	2.4.0	SB	04/21/08

Table of Contents

Table of Figures and Tables

PRODUCT BRIEF:
Logic embedded product solutions

LOGICLOADER™ Bootloader/Monitor

LogicLoader™ is a bootloader/monitor program developed by Logic Product Development that initializes an embedded device and is capable of loading both operating systems and applications. In addition, LogicLoader provides a full suite of commands for hardware configuration, in-field device management, hardware debug, manufacturing, and test.

Customizable and extendable at the user level, LogicLoader is built for multiple processor platforms (ARM, ColdFire, i.MX, SH, XScale), with support for both CompactFlash FAT and YAFFS file systems. LogicLoader contains a fully integrated TCP/IP stack—with DHCP and TFTP support—providing network bootstrap support. Greater customization to your specific needs can be achieved through conditional scripting and the ability for LogicLoader to drive LCD displays to show custom splash screens, making LogicLoader an excellent tool to fast forward your embedded product design.

Product Features

Operating System (OS) Bootstrap

+ Load multiple OSes (Microsoft Windows Embedded CE, Linux, etc.)
+ Load an OS from CompactFlash, resident flash array, serial connection, or Ethernet connection
+ Fully configure a hardware platform for the OS
+ Activate custom software functions to initialize hardware before the OS starts
+ Power-on self test capability

In-field Device Management

+ Modify boot actions at run-time
+ Remote device management eases debugging and upgrading

Hardware Debug

+ Link in custom test functions to verify custom hardware
+ Use a familiar UNIX-like interface for debugging the device
+ Ethernet-based download and debug interface for Windows Embedded CE

Custom Applications

+ Use LogicLoader to load, burn, and jump to any custom embedded application

Manufacturing and Test

+ Add in custom functional test software for your specific device needs
+ Take advantage of the fast Ethernet connectivity to reduce manufacturing test time

Download Formats

+ SREC
+ ELF
+ BIN
+ RAW

LOGICLOADER™ :: HIGHLIGHTS:

+ Ships standard on all Logic System on Modules
+ Conditional scripting
+ YAFFS (flash file system)

LOGIC WEBSITE :: DESIGN RESOURCES:

+ Logic Technical Support : http://www.logicpd.com/support/
+ Technical Discussion Group : http://www.logicpd.com/support/tdg/
+ Frequently Asked Questions (FAQ) : http://www.logicpd.com/support/faq/
+ For more information contact Logic Sales : product.sales@logicpd.com

LOGIC
embedded product solutions
411 N. Washington Ave. Suite 400
Minneapolis, MN 55401
T : 612.672.9495 F : 612.672.9489
I : www.logicpd.com

1.2 Acronyms

API	Application Programming Interface
BIN	Microsoft BIN file format
CPLD	Complex Programmable Logic Device
CF	CompactFlash®
DHCP	Dynamic Host Configuration Protocol
EEPROM	Electrically Erasable Programmable Read-Only Memory
ELF	Executable Linkable Format
FAT	File Allocation Table
FATFS	File Allocation Table File System
GPIO	General Purpose Input Output
GNU	GNU is not UNIX
IO	Input/Output
IP	Internet Protocol
JTAG	Joint Test Action Group
LAN	Local Area Network
LwIP	Lightweight implementation of the TCP/IP protocol stack
OS	Operating System
RAM	Random Access Memory
RAW	RAW file format, e.g., absolute binary
RISC	Reduced Instruction Set Computer
SOC	System on Chip
SOM	System on Module
SRAM	Static Random Access Memory
SREC	Motorola S-Record file format
TCP/IP	Transport Control Protocol/Internet Protocol
TFTP	Trivial File Transfer Protocol
YAFFS	Yet Another Flash File System

1.3 Technical Specifications

Please refer to the component specifications and data sheets applicable to your SOM:

- SOM Hardware Specification
- Applicable Processor Manual

1.4 LogicLoader Command Description Manual

For a complete description of LogicLoader's 'losh' commands, please see the *LogicLoader Command Description Manual* available from Logic's downloads page http://www.logicpd.com/auth/. The *LogicLoader Command Description Manual* explains how to use each LogicLoader command.

1.5 LogicLoader Addendums

Logic has written a SOM-specific addendum for each SOM that runs LogicLoader. LogicLoader Addendums are located under the "User Manuals" heading on Logic's Registered Products downloads page.

1.6 LogicLoader Labs

Logic has written informal labs that provide a step-by-step introduction to basic LogicLoader commands and usage for specific SOM platforms. These labs are available for download under the "User Manuals" heading on Logic's Registered Products downloads page.

2 LogicLoader

2.1 LogicLoader Overview

The LogicLoader (LoLo) is a bootloader/firmware-monitor program developed by Logic Product Development. LogicLoader is designed to initialize an embedded device, load and bootstrap an operating system, and provide a low-level firmware monitor with debugging functionality.

2.2 LogicLoader Basics

Most operating systems rely on an underlying bootloader to initialize a device from its reset condition. In general, operating systems are designed with the assumption that the system will be in a specific pre-defined state before the operating system is started. Some example assumptions might be that system RAM has been initialized and cleared, processor interrupts are disabled, and a timer has been initialized to provide a system tick for the OS. The LogicLoader program initializes Logic Product Development's SOM platforms and prepares them for use by an operating system.

Another basic function of LogicLoader is the capability to upgrade device software (flash memory, CPLD firmware, serial EEPROM contents) after deployment. This "in-field upgrade" ability requires a bootloader program that is capable of loading software images from various sources, as well as committing loaded images to non-volatile memory. LogicLoader implements this by giving the system the ability to load system software from flash memory, a CompactFlash storage card, a Local Area Network, or from a device attached to the system's serial port. LogicLoader also has the ability to upgrade an existing operating system residing in system flash.

LogicLoader was developed to fulfill the need for an OS- and processor-independent bootloader that can interface with a variety of hardware transports. The GNU development tool chain used to build LogicLoader is cross-platform capable.

2.3 Using LogicLoader for Debugging

LogicLoader implements a feature-rich firmware monitor, including the LogicLoader shell, also known as "losh." Losh is a command interpreter providing control over system state prior to loading an OS image. It has features such as command recall, command line editing, automated control via scripting, and diagnostic routines.

Losh includes many commands designed specifically to help software and hardware engineers debug low-level interfaces. For example, formatted data in arbitrary memory locations can be read from and written to using the 'x' and 'w' commands. Other commands run specific tests designed to verify Logic's SOM hardware platforms. All commands return a value to the command line that can be used to conditionally evaluate the command result. Refer to the *LogicLoader Command Description Manual* for a complete description of available commands.

Developers may code their own test programs using the provided GNU development tool chain and use the LogicLoader to load and run their software. This provides the ability to verify and debug hardware interfaces without the overhead of building, downloading, and running large operating system images.

2.4 Manufacturing Advantages with LogicLoader

LogicLoader can be used with a desktop software utility to load a device's system software on the manufacturing line. This utility is customizable to suit your desired transfer mechanism and additional needs. LogicLoader can also be augmented with functional test software to completely verify a device before it leaves the manufacturing line. Here is an example scenario: LogicLoader launches a device's final functional test at the end of a manufacturing line, and then loads the device's final software image before packaging. Contact Logic for more information on using LogicLoader to streamline manufacturing.

3 The LogicLoader Shell (losh)

3.1 Losh Overview

Losh is a command interpreter similar to those found in Unix environments. Losh implements a rudimentary network and file system command set, enhanced with custom diagnostic and memory manipulation commands for debugging hardware.

Developers familiar with a Unix-like command line interface should find the losh implementation familiar and easy to work with. Many of losh's commands are patterned after their Unix counterparts and share the same syntax.

3.2 Losh Basics

Losh uses a standard output stream (stdout). By default, stdout refers to a SOM's debug serial port. The output of any command that displays information to stdout (e.g., the 'cat' command) can be viewed using the terminal emulation program connected to the SOM's debug serial port. Likewise, the standard input stream (stdin) by default also refers to the SOM's debug serial port.

The LogicLoader shell includes a virtual file system that uses standard Unix path names. The highest-level (or root) directory is designated by the identifier '/'. A special sub-directory of the root with the name "dev" is used to enumerate and interact with the system's various peripherals and their associated device drivers.

3.2.1 Using Losh

The losh shell includes a basic command line editing feature and a command history feature. This provides you with a quick way to repeat commands. Using the up and down arrow keys, you can scroll through the list of previously executed commands. When a desired command is displayed, press the return key to repeat the command. The right and left arrow keys move the cursor anywhere within the current line. This allows you to modify, delete, or insert text anywhere in the current line without having to "backspace" the entire line and re-type commands.

Losh includes a user help feature through the 'help' command. Typing 'help' followed by any command name at the losh prompt will display the command's syntax, usage, and an example. This may be especially helpful to users who are just becoming familiar with the LogicLoader shell.

Commands may be run in the background by adding an '&' suffix.

3.2.1.1 Understanding the 'ls' Command

The 'ls' command lists the contents of the current directory. A sample terminal output that results from running the 'ls' command is shown below:

```
losh> ls
    :          NK.BIN    4268863
D   :             DOC          0
D   :            BOOT          0
```

In this example, the columns displayed are (in order from left to right): entity attribute, entity name, and entity size. See Figure 3.1, below.

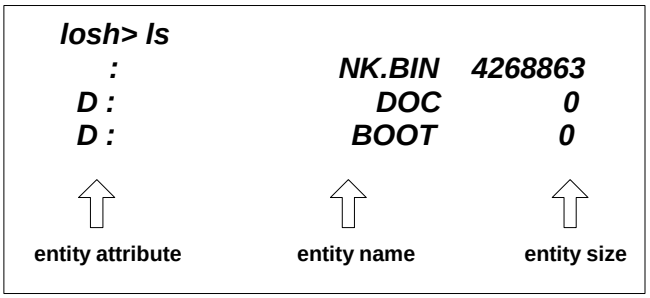

Figure 3.1: 'ls' Command Columns

The first column, entity attribute, can be blank, "D", "S", "R", "r" or "H". A blank field indicates a normal attribute, a "D" indicates a directory attribute, an "S" indicates a device driver attribute, an "R" indicates a read-only attribute, an "r" indicates that reserved bits are set, and an "H" indicates a hidden attribute.

The second column, entity name, is the name of the entity as it exists on the file system. This name should be used, with attention to case, in any commands referencing the entity.

The third column, entity size, indicates the size (in bytes) of the entity on the storage device.

4 Flash Devices and LogicLoader

LogicLoader supports both NOR and NAND flash devices; however, the usage is entirely dependent upon the available flash type(s) on your SOM (i.e., some Logic SOMs only have NAND, some only have NOR, and some have both NAND and NOR). NOR flash devices are linear, memory-mapped devices that can be read in a similar manner to any random access memory (RAM) device. Programming NOR devices requires a programming algorithm. LogicLoader supports NOR flash devices conforming to the Common Flash Interface (CFI) specification, which includes most NOR flash devices used today. NAND flash devices are block devices that require read and write algorithms. As of the time of writing this document, there is no common algorithm used to read or write to NAND flash devices; every manufacturer requires a unique algorithm.

4.1 NOR Addressing

Reading from a NOR device occurs in a similar manner to any RAM device. Writing to NOR devices, however, is a little more complicated. The default state for NOR flash is for each bit to be set at "1". Halfwords can be used to set bits from "1" to "0"; however, writing to a NOR device can only set bits from "1" to "0". In order to set a bit from "0" to "1", the entire block containing that bit has to be erased (i.e., all bits in that block are returned to their default state of "1").

Despite the similar addressing scheme between NOR flash and RAM devices, NOR flash cannot be used as a RAM device because NOR is block-organized to allow for erasing. The fact that NOR can be read as RAM is only used at boot-time, when it can be used as a permanent byte addressed storage device. When NOR is used as a file system device, block addressing is used.

4.2 Booting from NOR

LogicLoader resides in NOR flash starting at block 0. At reset, LogicLoader relocates itself from flash memory to system SDRAM and then spends the remainder of its run-time executing out of system SDRAM.

4.3 NAND Addressing

NAND devices use an addressing scheme of block, page, and sector. A block is the smallest erasable chunk of memory, whereas pages and sectors are merely mechanisms that describe the addressing hierarchy (blocks are made up of pages; pages are made up of sectors). The number of blocks, pages, and sectors will be unique for each particular NAND flash device. Some NAND devices may not have any sectors, in which case addressing is performed using only block and page.

NAND devices currently come in two flavors where addressing is concerned: small page and large page. Small page devices have a page size of 512 bytes; large page NAND devices have a page size of 2048 bytes. Larger page sizes tend to offer higher densities of NAND flash.

Whether the smallest chunk of data is addressed using a page or a sector, there is a spare area associated with that smallest chunk. This spare area will be 16 bytes for small page type devices and 64 bytes for large page devices. The spare area is used by software to manage:

- Error correction codes to correct single bit errors and to identify two or more bit errors.
- Manufacturer bad block identification.
- Flash file system metadata. The specific metadata will be unique to the particular flash file system used. LogicLoader dedicates a portion of the NAND spare area to YAFFS.

4.4 NAND Bad Blocks

NAND devices can develop bad blocks over time, as well as contain bad blocks when shipped from the manufacture. Bad blocks are defined as having two or more bit errors within the block. Single bit errors need to be corrected with software using an ECC algorithm. Most NAND blocks can be erased and rewritten on the order of 100,000 cycles before potentially going bad. NAND manufacturers state that the device integrity decays only with erase/program cycles. However, some third-party studies indicate that data integrity may decay with a large number of read cycles as well. LogicLoader and YAFFS assume data integrity does not decay with reads. YAFFS assumes writes may lose integrity over time, so NAND writes are all verified and two or more bit errors will result in YAFFS marking the block bad.

Unlike NAND flash, NOR flash devices do not have bad blocks.

4.5 NAND Programming

NAND devices are programmed by sending commands to the device. Similar to NOR devices, programming NAND consists of an erase phase that fills the entire block with 1s and a program phase that writes 0s to the device. Since NAND is a block device, a flash file system is needed to manage where data is read from and written to in order to avoid bad blocks on the device.

4.5.1 Skip Bad Block Method

A common algorithm used to program flash devices in production is the "skip bad block method." This is a flash file system in its simplest form. As the name implies, data is written contiguously on the device from low numbered blocks to higher numbered blocks, while skipping any bad blocks (as marked by the manufacturer). This algorithm works well for programming a NAND device once, but is not capable of removing and rewriting portions of the written image.

4.5.2 YAFFS Overview

The YAFFS file system has been optimized for NAND use. YAFFS is able to:

- Identify and avoid bad blocks using an ECC algorithm.
- Use load leveling, where erasing and writing is averaged out among all the blocks of the device, and no one block is erased and written repeatedly.
- Manage metadata, such as directories and links.

More information regarding how YAFFS operates in LogicLoader can be found in the "YAFFS" Section of this document.

5 Block Devices

Within LogicLoader, a "block device" is any device that only contains pages and sectors that can be read from or written to (this includes devices that require a block erase before a write). Examples of block devices are: ATA, NOR flash, and NAND flash. The LogicLoader shell supports all block devices with the same command set. More detailed information for each command can be found in the *LogicLoader Command Description Manual*.

5.1 Using Block Reference

In the losh command set, there are two different methods to directly reference a block on a device. Some commands require a "B" to be placed in front of the block number. For example:

```
losh> erase/dev/nand0 B9 B500
```

In this example, omitting the "B" would indicate flat memory addressing. Use of flat memory addressing is discouraged and should be replaced with block-aligned memory addressing.

Other commands require the block address to be written without a "B" in front of the block number. For example:

```
losh> part-add /dev/nand0 a 1 1024
```

This command creates a partition "a" in the NAND device. The partition starts at block 1 and is 1024 blocks long. (Partitions are discussed in detail in Section 10 of this document.)

Because the specific command dictates the proper method to reference a block, it is important to understand the requirements of that specific command. The 'help' feature may be useful in determining which method should be used.

5.2 burn

The 'burn' command works with any block device. It burns the loaded image to the device from a given block offset. For example:

```
losh> burn /dev/flash0 5
```

This command will burn the loaded image to flash starting at block 5.

For burning to a NAND device with the skip bad algorithm, use the 'dd' command as described below.

5.3 dd

The 'dd' command copies blocks from a source device to a destination device. The command can use the skip bad block algorithm and it can be turned on with a flag. Use of the 'dd' command is not limited to block devices; it can be used with whatever device you want. However, the device used with the command does determine how to specify bock size. For NOR flash, the 'dd' command requires any halfword aligned block size; but for ATA or NAND flash the command requires you to specify the exact read/write page size. For ATA, the page size is 512; for NAND flash, the page size is either 512 or 2048. The 'info mem' command can be used to determine the correct page size for your NAND device.

For example:

```
dd if:/load of:/dev/nand0 count:16 ibs:512 obs:512 os:16 skip_bad:1
```


This example command will copy the contents of "load" into the NAND data area, skip spare blocks, and use the skip bad block algorithm. For command argument details, please see the *LogicLoader Command Description Manual*.

5.4 erase

The 'erase' command can be used with any device. However, extra caution must be taken when erasing NAND and NOR blocks so as not to erase a YAFFS partition or any LogicLoader files. An attempt to erase these files or partitions will require confirmation before the erase command continues; this will prevent mistakenly erasing files required by the system to boot. NAND blocks with bad block markers will not be erased. For example:

```
erase /dev/nand0 B10 B502
```

This command will erase 502 blocks of the device starting at block 10.

There is an optional argument "force" that can be used with the 'erase' command. The force argument will force the 'erase' command to erase all blocks in the specified address range even if they have been marked bad. Without the "force" argument, the 'erase' command will skip bad blocks in an effort to preserve bad block information. Extreme caution must be used when using the "force" argument. If a block has been marked bad by the NAND manufacturer, and the block is erased with the "force" argument, there is no way to ever recover the bad block information. For example:

```
erase /dev/nand0 B10 B502 force
```

Note: The legacy syntax `erase <offset> <length> <device>` is only supported for compatibility reasons.

5.5 info

The 'info' command can be used to return specific information about the NAND and NOR devices, as well as information about any YAFFS boot partitions. This information is returned by using the 'info mem' and 'info YAFFS' arguments. The 'info mem' command includes geometry data for NAND and NOR flash devices. The geometry information includes:

- Base address (unique to NOR devices)
- Number of blocks
- Bytes per block
- Is chunk device (if 'is_chunk_device' equals 0, then the following information is not relevant and, therefore, is not printed)
- Number of chunks
- Chunk size
- Bad block list
- Bytes per chunk
- Bytes per spare

5.6 update

The 'update' command is used to load and install an update image; it also includes support for updating LogicLoader in the YAFFS partition. When update files are sent to the SOM using the 'update' command, LogicLoader will identify the update as LogicLoader, and then program the NAND part as needed.

6 Program Loading

Using LogicLoader to download any application, operating system, or update to a device requires an understanding of the interaction between the 'load', 'burn', 'jump', and 'exec' commands. The purpose of this section is to describe each individual command and explain the interaction between these commands.

6.1 Understanding the 'load' Command

The purpose of the 'load' command is to transfer an executable image to a device. The image must be in one of the following supported formats: ELF, SREC, RAW, or BIN. The 'load' command uses information inherent to the supported formats (or as entered as part of the command for RAW format) to determine where the downloaded image should be stored in the device's memory. The 'load' command stores the destination address of the downloaded image for later use by the 'burn' command, and stores the program start address for later use by the 'jump' or 'exec' commands. For RAW format, 'load' will store the destination address as the program start address. The image must be destined to reside in either flash memory, system RAM, or on-chip SRAM.

If an image is destined for system RAM or on-chip SRAM, the 'load' command stores the image directly to its run-time location. Refer to Figure 6.1 for a graphic representation of this process.

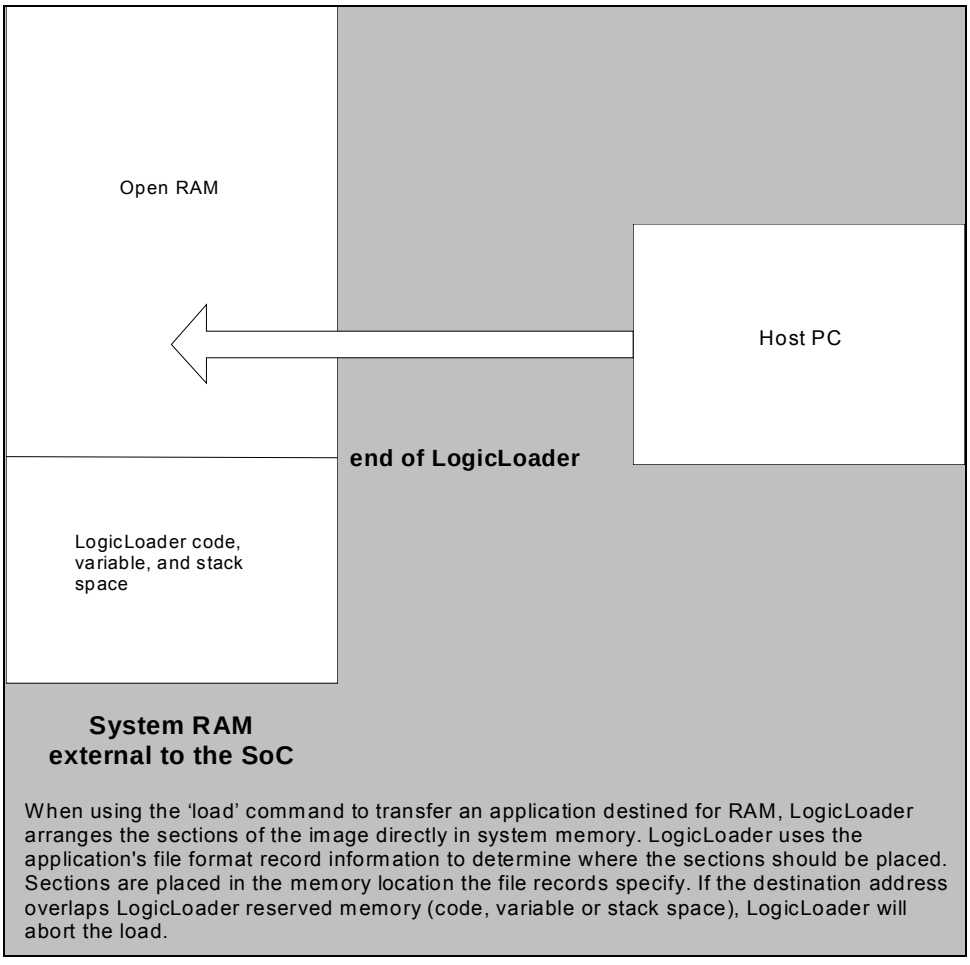

Figure 6.1: Downloading to RAM

If a downloaded application is destined for flash memory, the 'load' command transfers the file into a temporary RAM buffer on the device. The transferred image may be programmed into flash

using the 'burn' command after the transfer is complete. Refer to Figure 6.2 for a graphic representation of this process.

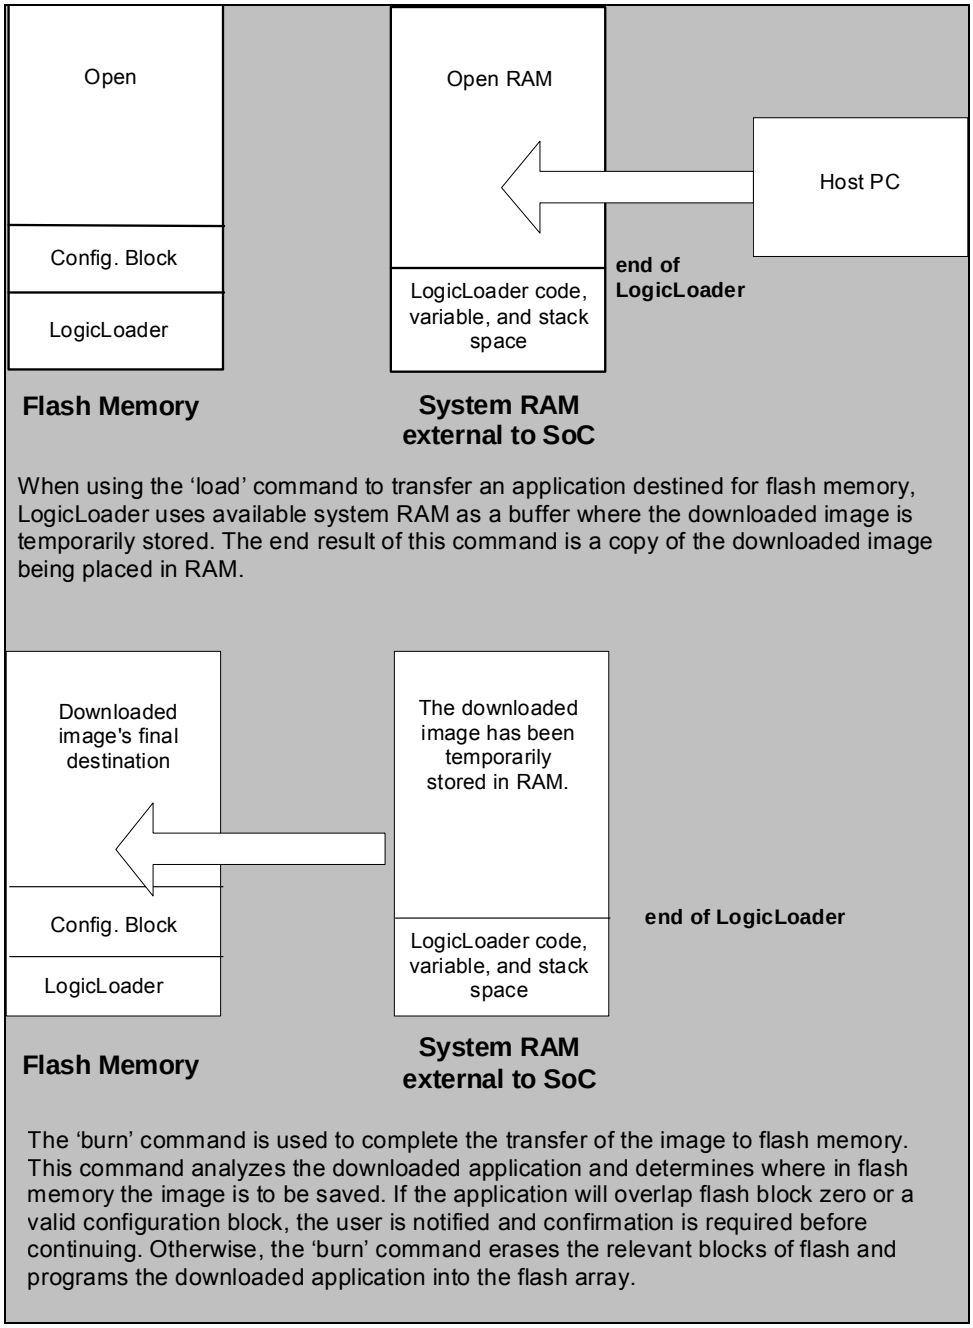

Figure 6.2: Downloading to Flash

6.1.1 Using TFTP as a Source

A file located on a TFTP server can be used as the source for the following commands: 'load', 'cat', 'hd', 'md5sum', and 'cp'.

The general form for a TFTP file is "/tftp/<server>:<filename>:[port]" where <server> is the IP address of the server, <filename> is the name of the file on the TFTP server (including subdirectory identifiers), and [port] is the optional port number the TFTP server is listening to. If nothing is specified for the port, it is assumed the TFTP server is using the standard port 69.

For example, to load the ELF file "image.elf" from a TFTP server accessible at IP address 192.168.3.6 that is listening on the standard port, the following command would be used:

```
losh> load elf /tftp/192.168.3.6:data-file
```

Another example would be to load the Platform Builder file "NK.bin" from the TFTP server at IP address 10.1.240.10 listening on port 3001:

```
losh> load bin /tftp/10.1.240.10:NK.bin:3001
```

6.2 Understanding the 'burn' Command

The 'burn' command should only be used following the successful download of a binary image destined for flash. If the 'load' command is used to download a flash image, the image is temporarily stored in a reserved section of system RAM. The 'burn' command is responsible for actually erasing the necessary blocks and programming the downloaded image into flash at the destination address. Refer to Figure 6.2 for more information.

6.3 Understanding the 'jump' and 'exec' Commands

LogicLoader provides two different ways to transfer execution to your application. The 'jump' command is more useful for launching and debugging an application that will be relying on LogicLoader or an operating system to setup the run-time environment. The 'exec' command is more useful for launching an application, such as an operating system that will take over total control of the hardware and the environment. The differences between the 'jump' and 'exec' command are that only 'exec' can pass a command line argument to the program being executed and that 'exec' disables interrupts, the cache, and the MMU (if present).

6.3.1 The 'jump' Command

The 'jump' command is an assembly-level jump to the starting instruction of a program. If 'jump' is executed without a parameter, LogicLoader will jump to the program start address of the last program loaded to system RAM (if any). If an address is passed in, the 'jump' command will jump to the specified address. After a 'jump' command is performed, LogicLoader continues to execute in the background. LogicLoader does not set up a run-time environment for a program, rather the program inherits LogicLoader's current environment. It is the software engineer's responsibility to ensure that the hardware is setup in the desired manner.

This example may be used when writing a function that LogicLoader will 'jump' to:

```
int my_jump_function(void);
```

6.3.2 The 'exec' Command

The 'exec' command is an assembly-level jump to the starting instruction of a program that will pass in three arguments. If 'exec' is executed without a parameter, LogicLoader will jump to the program start address of the last program loaded to system RAM (if any) and pass in a pointer to an empty string. If both an address and command line are specified, the 'exec' command will jump to the specified address and pass a pointer to the command line provided. The 'exec' command will disable interrupts, the cache, and the MMU (if present) prior to executing the jump.

The 'exec' command passes the command line argument via a pointer to memory that has been allocated from LogicLoader's heap. Any application or OS code must preserve the command line, or finish using the command line arguments, before reclaiming LogicLoader's memory space for its own use. Because the 'exec' command shuts off the MMU, the image must have a virtual address that maps directly to its physical address since the entry address that 'exec' jumps to will always be a physical address.

This example may be used when writing a function that LogicLoader will 'exec' to:


```
int my_exec_function(unsigned int arg1, unsigned int arg2, char *cmd_string);
```

The first two arguments, arg1 and arg2, have different values depending on the flags given to 'exec'. The third argument will be a pointer to the command line, as described above.

To boot an ARM Linux kernel, use the '-t' argument with the 'exec' command. This causes arg1 to get zero, arg2 is then the architecture ID, and arg3 is a pointer to an ATAG structure that contains, among other things, a pointer to the cmd_string.

6.3.3 Command Example Using 'load' and 'burn' with 'jump' or 'exec'

An application program that is written for the Zoom Development Kit can be linked to reside in flash or ram.

First, let's assume that we have built an application for flash. To properly store this program in flash, issue the 'load' command followed by the 'burn' command. Make note of the program start address (for example: 0x400d0100) so that you can jump to the program after a reset. Once the image has been burned to flash, you may enter the 'jump' or 'exec' command specifying 0x400d0100 as the argument at anytime but you can take a shortcut if you have not reset the board since the 'load' command will store the program start address. A valid sequence would be as follows:

1. `losh> load elf`
   This transfers the image to the device.
2. `losh> burn`
   This programs the image into flash at the destination address stored by the 'load' command.
3. `losh> jump or exec`
   This will work because the 'load' command saved the program's flash start address. Both the burn destination address and the program start address will be valid until the next reset or the next use of the 'load' command.

After a reset the program may be launched at any time using the 'jump' or 'exec' commands with a specific destination address:

```
losh> jump 0x400d0100
```

or

```
losh> exec 0x400d0100 –
```

Next, let's assume that we have built an application for RAM. To properly load and execute an application out of RAM, issue the 'load' command followed by the 'jump' or 'exec' command. A valid sequence would be as follows:

1. `losh> load elf`
   This transfers the image to the device.
2. `losh> jump or exec`
   This will work because the 'load' command stored the program start address. The program start address will be valid for this program until the next reset, or the next use of the 'load' command.

Keep in mind that the option of specifying the program start address, as shown in the flash example, is also available.

6.4 Understanding the 'update' Command

Logic deploys software or firmware updates in the form of update files (.upd extension). To deploy an update file, use the 'update' command. If a filename/path parameter is not passed to the 'update' command, the system will assume that stdin is being used to send the update file to the

system. When the update command is activated, after the system has received the .upd file, it automatically launches the file and performs the actions required.

Update files are comprised of self-extracting applications that, once activated by the update command, run and perform whatever function the application was coded to carry out. This allows a single "update" command to perform a variety of different actions from a self-contained file with minimal user interaction.

The procedure to update LogicLoader with the 'update' command differs from the 'load/burn' procedure in this way: only one command implements the entire update process without any user interaction or confirmation.

7 Scripting

7.1 Scripting Overview

Scripts can be used to automate any commands or command sequences entered on the command line. Scripts are comprised of a simple text file with a listing of commands that the user wishes to automatically execute in sequence.

7.1.1 Scripting Rules

Basic scripting rules are as follows:

- Enter commands into the script file with the same syntax used on the losh command line;
- Separate commands with a semi-colon or a new line;
- End the script with a '\n' (this tells the parser to stop parsing the file and instructs the command interpreter to start executing the script);
- Use the command 'exit' to end the script (this tells the command interpreter to stop executing the script).

7.2 Launching Scripts

The process of launching a script manually, or post-boot time, uses the 'source' command. For example: the command `source /cf_card/myscript.txt` will execute the script stored in the file "myscript.txt" on a mounted CompactFlash card. For more information on the 'source' command, please refer to the *LogicLoader Command Description Manual* document.

The process of auto-launching scripts on startup is referred to as "boot-time scripting." Boot-time scripts are the primary mechanism used for automatically launching an OS or application when deploying a product to the field. Their capability is the same as other scripts, with the difference being their ability to be automatically run at startup. You can think of a boot-time script fulfilling the same role as an "autoexec.bat" file commonly found on desktop operating systems. Boot-time script usage is described more thoroughly below.

A third way to launch a script is to "send" it to the system while LogicLoader is waiting at the losh prompt. If the script file is sent over the terminal emulator connection to the losh shell, the script will be entered on the command line as if typed in by the user. If the script being sent incorporates a carriage return at the end of the script, the command line will launch the script when it receives the carriage return. This type of script launching is primarily used during development when the developer wishes to send a number of development commands to LogicLoader in sequence. For example: a command sequence initializes the Ethernet interface, downloads a Windows CE OS image, and then launches the OS image with a specific command line.

7.3 Persistent Script Storage

In order for a script to persist across power cycles, the script must be stored in a local, non-volatile memory device on the system. There are a number of different persistent storage locations that can be used to store a script. The primary storage mechanisms supported by LogicLoader are the serial EEPROM, the resident flash array (dev/config or YAFFS), and the CompactFlash interface. Because different SOMs may not have one or more of these interfaces available in hardware, please refer to the individual SOM's *LogicLoader User's Manual Addendum* document for specific persistent storage interface support.

7.3.1 Persisting Scripts with the Echo Command

The 'echo' command can be used to store a script in the serial EEPROM or /dev/config. To include a new line in the first argument to 'echo', it is necessary to enclose the whole argument in

double-quotes. Remember to end the script by inserting '\n' before the end quotes to instruct the parser to stop parsing the file. Since scripts stored in the serial EEPROM or /dev/config are not stored as actual files, it is important that any previous information in the serial EEPROM or /dev/config is not interpreted as part of the script. Check the contents of the serial EEPROM or /dev/config with 'cat' or 'hd' to verify that the contents are as expected. If not, the 'erase' command should be used to erase any previous information before the 'echo' command is executed.

7.3.2 Serial EEPROM Scripts

The system's serial EEPROM is one persistent storage area that supports the storage and execution of scripts. The serial EEPROM is the primary boot-time script storage location. Boot-time scripts stored to the serial EEPROM are typically short and may redirect to a secondary script on an interface capable of larger storage capacity.

To store a script to the serial EEPROM interface (/dev/serial_eeprom), use the 'echo' command. An example of using the 'echo' command to store information to the serial EEPROM is shown below:

```
echo "LOLOmount fatfs /cf; source /cf/B.BAT; exit;\n" /dev/serial_eeprom
```

7.3.3 Configuration Block Scripts

The system's configuration block is another persistent storage area that supports the storage and execution of scripts. The configuration block is located in system flash memory and is the secondary boot-time script storage location on systems with serial EEPROM.

Store a script to the configuration block interface (/dev/config) by using the 'echo' command or the 'config S' command. Here is an example of using the 'echo' command to store a script to the configuration block:

```
echo "LOLOmount fatfs /cf; source /cf/B.BAT; exit;\n" /dev/config
```

Note: The configuration block must be initialized before using it for scripting commands. For more information on how to create and use the configuration block, please see Section 9 of this document.

7.4 Settings that Affect Scripts

The 'set' command can be used to modify several internal variables affecting script execution. These function similarly to the Unix shell scripting analog, where a '-' causes the flags that follow to be set, and a '+' causes them to be unset. It is highly recommended during development to set the '-w' flag to receive warnings about common scripting errors.

The flags available are:

- e Exit script execution immediately when commands fail;
- n Read commands, but do not execute; ignored by interactive shells;
- q Do not print LogicLoader error messages;
- u Exit on expansion of unset variables;
- v Echo input lines as they are read;
- w Print warnings for possible errors;
- x Echo all user commands before executing them.

7.5 Using Boot-time Scripting

It is possible to execute a script automatically at startup. This is useful for making the device jump into an operating system or other program when powered-on without requiring manual command-line input. This functionality can be described as being equivalent to the system automatically calling the 'source' command on one of the boot-capable devices.

7.5.1 Boot-time Script Guidelines

All of the commands available in LogicLoader are also available to boot-time scripts. As in normal scripts, a semi-colon must be used to separate commands and the exit command must be used to terminate a boot-time script.

In order for a script to be boot-capable, the script must be stored in a boot-capable location and must contain the necessary "magic" string prefix. A boot-time script may reside either in the on-board serial EEPROM or in the flash-based configuration block. The order of boot-time execution is first the EEPROM, then the configuration block. In order to differentiate between auto-booting scripts and non auto-booting, LogicLoader checks the first four bytes of the boot-capable devices to see if they contain a "magic" string indicating that the following script should run automatically.

7.5.2 Boot Script Magic Strings

The "magic" string for LogicLoader is "LOLO" for silent execution or "VOLO" for verbose script execution. If the string "LOLO" prefixes the boot script, the script's commands and terminal output will be completely silent. By using "LOLO" as the prefix, it is possible to fully boot the system without ever sending any information out the debug serial port. If "VOLO" is used as the boot script prefix, the boot script commands, return codes, and other "normal" information is displayed via the serial port as if the script was running post-boot time.

7.5.3 Exiting a Boot Script

A common need is to abort the execution of a boot script in order to exit into the command line for additional debugging, development, or simply to change the boot script. The primary way to accomplish this is by holding the 'q' key down in a terminal emulator program attached to the device's debug serial port.

The system does pause for one/half of a second to read from debug serial port to determine if an abort request is being made. Some of Logic's SOM products implement an external mode line that allows LogicLoader to ignore the assertion of the 'q' key – thereby skipping the one/half second wait time and decreasing the overall boot time of the system when a boot script is desired. For more information on the hardware line that provides this functionality, check the *LogicLoader User's Manual Addendum* for your hardware

7.5.4 Understanding the Echo Command

The 'echo' command can be used to store a script in the serial EEPROM or /dev/config. The 'echo' command only writes the number of bytes contained in the string. If the string to be written is shorter than the previous contents, the result of the echo will not be what is intended. Use the 'cat' or 'hd' command to verify the contents of the serial EEPROM or /dev/config before using the 'echo' command.

7.5.5 Boot-time Script Example

The following example creates a simple LogicLoader boot script that first mounts the CompactFlash card and then runs a second script "B.BAT" on the CompactFlash card that automates software.

```
losh>echo "LOLOmount fatfs /cf; source /cf/B.BAT; exit;\n" /dev/serial_eeprom
```

or

```
losh>echo "LOLOmount fatfs /cf; source /cf/B.BAT; exit;\n" /dev/config
```


7.6 Conditional Scripting and Variables

7.6.1 Variables

LogicLoader's shell supports the concept of shell variables. The syntax and usage of these variables are patterned after the BASH shell.

7.6.1.1 Variable Names

A variable name may be any sequence of letters, numbers, or the underscore token.

7.6.1.2 Variable Assignment

A variable is created and assigned a value by using the '=' operator. For example:

```
losh> foo = 1
```

creates a new variable named "foo" and assigns it the value of "1". Once a variable has been created, it may be assigned a new value at any time by using the '=' operator again. The right-hand side of an assignment statement is not limited to a simple number; it can be a complex expression involving other variables.

7.6.1.3 Internal Representation

Variables are internally represented as strings. For example:

```
losh> foo = 1
```

internally points the variable "foo" at a sequence of characters equivalent to: 0x31 0x00. Because variables are treated as strings, commands may be aliased as variables. For example:

```
losh> e = echo
losh> msg = "Hello World"
losh> $e $msg
Hello World
```

Notice the quotes used to ignore white space. If the created variable will be assigned to more than one token, the tokens must be included in double-quotation marks.

7.6.1.4 De-referencing a Variable

To dereference a variable, that is, to access a variable's assigned value, use the '$' operator. For example:

```
losh> foo = "Hello World"
losh> echo $foo
Hello World
```

The '$' operator causes the shell to substitute the variable with the string value assigned to it. In some cases, a variable's assigned value will be converted into a numeric value. This occurs when the shell is evaluating a conditional expression. This is described in more detail below.

If a variable is referenced that does not have a previous value, its value is assumed to be zero and a warning message is printed.

Note: Enclosing a sequence of tokens within double-quotes binds them together into a single token. For example:

```
losh> e = "echo Hello World"
losh> $e
echo Hello World: command not found
```


will not work because the parser only evaluates the string once. Thus, instead of being split up into three distinct tokens, the double-quotes cause the tokens to be bound and treated as one.

7.6.1.5 Built-in Variables

The shell contains two built-in variables, '?' and '@'.

The '?' variable is assigned to the return value of the last command executed. By convention, all shell commands return zero to indicate that it completed successfully and a non-zero error code to indicate a failure. To view a command's return value, use the 'echo' command and the value of the '?' variable. For example:

```
losh> mount fatfs /cf   # Mount a FAT file system.
losh> echo $?           # Display the value returned from the mount command.
```

The '@' variable is an auxiliary variable that is set by some commands. For instance, the 'echo' command sets this value to the number of characters that it wrote. Therefore:

```
losh> echo "Hello"
losh> echo $@
0x5
```

The number "5" is printed because the string "Hello" contains five characters.

Please reference the *LogicLoader Command Description Manual* for specific command descriptions in order to learn which commands set the '@' variable, and if so, the usage of these commands.

7.6.1.6 Conditional Scripting

LogicLoader's shell supports an 'if-else-endif' programming construct as well as a 'while' construct. The syntax for an if-statement and an if-else statement is shown below:

```
if (expression)
      action
endif


if ( expression )
      action-1
else
      action-2
endif
```

Parentheses are not required around the expression, but they are encouraged to improve readability of the script. Similarly, tabs and new lines are not needed. The various elements of the construct may be separated by the ';' operator if so desired. For example:

```
losh> if expression echo "pass"; else echo "fail"; endif
pass
```

or

```
losh> if expression echo "pass"
      else echo "fail";
      endif
```

The syntax for a 'while' statement is shown below:

```
while ( expression )
```



```
        action
done
```

The expression is evaluated first. If the return is non-zero, then action is taken and control comes back to the expression evaluation. This is repeated until the expression evaluates to zero.

Note that 'if' and 'while' statements can be nested. The following example calculates the greatest common divisor of the numbers stored in the variables 'a' and 'b', leaving the result in 'a':

```
losh> while ($a .ne $b) {
            if ($a .gt $b) {
                  a = $a - $b
            } else {
                  b = $b - $a
            }
      done
```

7.6.1.7 Expressions

An expression is defined as a number or a combination of a logical operator and a number or numbers. If a variable has been defined and is being de-referenced in an expression, its value is converted to a number. An expression evaluates to true if the result is non-zero and false if the result evaluates to zero. Therefore, the simplest expressions would be:

```
if ( 1 )    # evaluates to true.
if ( 0 )    # evaluates to false.
```

7.6.1.8 Using Shell Variables

```
losh> foo=1
losh> bar=0x0
if ( $foo ) # evaluates to true.
if ( $bar ) # evaluates to false.
```

The other operators supported by the shell are listed below in order of decreasing precedence.

'-' '!' '~'	unary minus, logical not, arithmetic not
'*' '/' '%'	multiplication, division, modulus
'+' '-'	addition, subtraction
'<<' '>>'	left shift, right shift
'.lt' '.le' '.gt' '.ge'	less than, less than or equal, greater than, greater than or equal
'.eq' '.ne'	equality, inequality
'<' '>'	less than, greater than
'==' '!='	equality, inequality
'$((' '))'	immediate evaluation open, immediate evaluation close
'^'	bitwise exclusive or
'|'	bitwise or
'&'	bitwise and
'&&'	logical and
'||'	logical or

Note that the operators '==', '!=', '>', '<' apply to either strings or integers, but the evaluation is done as string comparisons. The operators '.eq', '.ne', '.lt', '.le', '.gt', '.ge' apply to either strings or integers, but each side of the expression must evaluate to numbers.

Immediate expression evaluation:

The immediate evaluation construct '$((' … '))' is used when a command needs an immediate value In this case the expression contained in '$((' … ')) is immediately evaluated and returned as a number. For example, the 'x' command can not take an expression as its operand:

```
losh> x /x 0x80200000 + 0x10 4
error: x: wrong number of arguments
```

Using the immediate evaluation construct '$((' … '))' gives:

```
losh> x /x $(( 0x80200000 + 0x10 )) 4
0x80200010  04001000 eb000000 fe000001 40ea0003  ……@
```

The following are all valid expressions that can be used as the right-hand side of an assignment, as an argument to a command (if enclosed in an immediate evaluation construct), or as the conditional expression in an 'if' or 'while' construct:

```
1 & 0            # evaluates to zero
1 | 0            # evaluates to one
0x01 ^ 0x02      # evaluates to 0x3
1 >= 2           # evaluates to zero
0 .ge 1          # evaluates to zero
1 + 3 * 5 ^ 7    # evaluates to 23 (reduces to 16 ^ 7)
```

As mentioned above, the shell exports two built-in variables. These are '?' and '@'. The variable '?' holds the return value of the last command executed. Therefore, constructs like the one below can prove to be very useful:

```
mount fatfs /cf
if ( $? )
      # Save current return values because 'echo' will overwrite them
      s_q = $?
      s_a = $@
      echo "Error, mount failed error codes: "
      echo $s_q
      echo $s_a
else
      echo "Mounted FAT file system at point '/cf'"
endif
```

Note: In the case of an error, the values of the '?' and '@' variables are saved. This is because the first call to the 'echo' command will overwrite the value of those variables.

7.6.1.9 Escaping the variable character

If the 'echo' command is used to store a variable reference in a script, the '\' operator must be used before that variable in order to defer evaluation of that variable until echoed. For example:

```
echo "if ($a == 2) source bar;\n" /dev/config
```

needs to be written as

```
echo "if (\$a == 2) source bar;\n" /dev/config
```

in order to prevent losh from evaluating the variable 'a' in the string before the echo call is used. This method applies to any string which must include a literal '$' character.

7.6.1.10 Comments

In order to make it easy to self-document scripts, the shell recognizes and ignores comments. A comment begins with the character '#' and extends to the end of the current line.

7.6.1.11 Numbers

The shell recognizes the following number formats:

- decimal
  - contains the characters 0-9
  - does not start with a zero
- octal
  - contains the characters 0-7
  - starts with a zero
- hexadecimal
  - contains the characters; 0-9, a-f, or A-F
  - starts with the sequence '0x' or '0X'

8 Video Interface

8.1 Video Interface Overview

LogicLoader includes the following video commands to configure the video controller:

- video-clear - clears the default video screen (sets the frame buffer to a monolithic color)
- video-close - turns off and un-initializes the default video device
- video-fb - sets or displays the video frame buffer address
- video-init - connects and initializes default video device settings, but does not enable the controller
- video-off - turns off an initialized display
- video-on - turns on an initialized display
- video-open - connects and initializes default video device settings and enables the display controller (equivalent of video-init and video-on)

8.2 Using the Video Interface after Initialization

Once the display has been initialized with either the 'video-open' or the video-init' commands, any of the drawing commands can be used. The 'video-fb' command allows the user to change the frame buffer address.

After executing the 'video-fb' command to change the frame buffer address, all drawing commands will use the new frame buffer address instead of the default. The 'video-init' command can be used to connect and initialize the video controller without enabling the video display. Then use the 'bitmap' command to draw to different areas in memory prior to using the 'video-on' command to turn on the display. A typical command sequence might look like the following:

```
losh> video-init 7 16
video-init display: width: 640 height: 480 bpp: 16 disp: 7
losh> bitmap TEST1.BMP 0xc0400000
losh> bitmap TEST2.BMP 0xc0600000
losh> video-fb 0xc0400000
losh> video-on
.....other command sequences
losh> video-fb 0xc0600000

.....other command sequences
losh> video-off
```

If using the configuration block, up to eight uniquely named custom screen settings can be saved. The stored settings include the frame buffer address so that the frame buffer will be initialized to a user specified address upon executing the 'video-open' or 'video-init' command. The current frame buffer address can be ascertained by issuing the 'video-fb' command.

9 Configuration Block

9.1 Configuration Block Overview

Logic has added an optional configuration block that is located in the first 64K of the second 256K block of flash, immediately following the location of LogicLoader. The purpose of the configuration block is to allow our customers the ability to store larger scripts, change the baud rate of the debug serial port, store default settings for the Ethernet, and save custom LCD controller settings. The script area accommodates 16kB of script storage space, and the peripheral settings allow for up to eight different settings each to be saved for the video, Ethernet, and serial. Custom settings can be downloaded and saved, or pre-programmed at the factory.

The configuration block is designed to be easily accessible from a customer's application. The structure and field definition header files are available from Logic. Each section has its own checksum, and there is also a checksum and version id for the entire configuration block. It is possible to store additional customer-defined settings to the configuration block for later use by customer software. In order to do this, a custom configuration block must be created. Please contact Logic for more information.

The configuration block is optional for most SOMs. Normal operation, with the default settings, is available to customers who do not wish to use the configuration block.

9.1.1 Initializing

The configuration block must be initialized on systems that have never implemented a configuration block. If, for some reason, the system's configuration block needs to be cleared, this initialization step will also provide for that.

Enter 'config CREATE' at the losh prompt to initialize (or re-initialize) the configuration block. The configuration block will be located in the 64K of flash immediately following the LogicLoader flash storage location.

9.1.2 Scripting

The default mode of the configuration block is scripting. Once the configuration block has been initialized, users can 'cat' and 'source' /dev/config in order to view or execute their script. The 'echo' command may be used to create the script, or larger scripts can be downloaded and saved using the 'config' command with the 'S' option. The scripts in the configuration block can be boot scripts or ordinary scripts. For more information on scripting, please refer to Section 7.

9.1.3 Video

A customer may save up to eight custom LCD screen settings, including the frame buffer location, by using the 'config' command with the 'V' option. In order to store a custom setting, set up the controller as desired and then enter 'config V screen_name X Y' to store it (where X and Y are dimensions of the screen—e.g., 640 and 480). Once this configuration step is implemented, the system can use the command sequence 'video-open screen_name depth" to use the new configuration.

9.1.4 Serial with the Configuration Block

A customer can change the baud rate of the debug port with the 'B' option. The setting will be stored under the debug serial port's UART index, and will automatically be used after the next reset. Only valid settings will be allowed.

9.1.5 Ethernet

A customer may save a default Ethernet setting that consists of a MAC address, and IP address, a subnet mask, and a gateway.

To accomplish this, set up the controller using the 'ifmac' and the 'ifconfig' commands, and then use the 'config' command 'E' option with an index (e.g., 'config E 0') to save the settings. The index used must correspond to the hardware name index used in the ifconfig command (e.g., sm0 – where 0 is the valid index).

Use the default setting by typing 'ifconfig sm0 /dev/config' (where sm0 indicates index 0) to re-load the values stored in the configuration block. In general, the MAC address is stored in the config block for viewing purposes only, and the actual MAC address used is accessed directly by the Ethernet chip out of its dedicated serial EEPROM. On SOMs without a serial_eeprom, the config block is required if the customer wishes to use the Ethernet feature from within LogicLoader.

The configuration block allows storage for up to eight Ethernet interface configurations, but the config E command is only able to store information for SOM hardware that is supported by LogicLoader.

10 Partitions

10.1 Partitions Overview

Theoretically, partitions can be created on every block device. However, the current implementation of LogicLoader only allows users to create partitions on NOR and NAND devices. Partitions on other devices, such as ATA, CompactFlash, and SD cards, can be used, but new partitions cannot be created on those devices. There can be up to four partitions on a device; however, extended partition tables are not supported.

Inodes are created for every partition at boot time. For example:

```
losh> ls /dev
```

will return the following:

```
losh> ls dev
        S  :                          sdmmc0d        0
        S  :                          sdmmc0c        0
        S  :                          sdmmc0b        0
        S  :                          sdmmc0a        0
        S  :                           sdmmc0        0
        S  :                            ata0d        0
        S  :                            ata0c        0
        S  :                            ata0b        0
        S  :                            ata0a        0
        S  :                             ata0        0
        S  :                           nand0d        0
        S  :                           nand0c        0
        S  :                           nand0b        0
        S  :                           nand0a        0
        S  :                            nand0        0
        S  :                          flash0d        0
        S  :                          flash0c        0
        S  :                          flash0b        0
        S  :                          flash0a        0
        S  :                           flash0  2097152
        S  :                             null        0
        S  :                            uart0        0
        S  :                           config        0
```

Within this example, nand0a is only an empty inode at this time and cannot be accessed (unless previous partition tables have been created on the NAND device). For example:

```
losh> hd /dev/nand0a 512
```

will return the following error message:

```
Partition does not exist, type 0xff
error: hd: failed to open (/dev/nand0a)
```

The most important thing to know about partition handling is that there is a RAM-based partition table for every device (referred to as RAM-partition table in this document). At boot, LogicLoader tries to fill this partition with data from the device. If it does not find a partition table on the device, then it will be empty. This means that it will be filled with 0s or 1s, depending on the type of device (the partition is filled with 1s for NOR and NAND flash devices; the partition is filled with 0s for CompactFlash and SD cards).

Returning to the example above, the nand0a RAM-partition table is filled with 1s. The partition driver reads the corresponding entry from the RAM-partition table and finds that its type is 0xff, or empty. This is why the error message states the partition does not exist.

So for every partition inode, there is a corresponding partition entry in the RAM-partition table.

10.2 Partition Creation in the RAM-Partition Table

Partitions can be created with the 'part-add' command. For example:

```
losh> part-add /dev/nand0 b 1 2048
```

will create a partition entry in the nand0 RAM-partition table. Note that the second (nand0b) partition of the device will be filled due to specifying b in the argument; this occurs because every device can have up to four partitions, which are labeled from a to d. The partition will be added beginning in block 1 and will have a length of 2048 blocks.

For example:

```
info part /dev/nand0
```

will output the table below. This command prints the RAM-partition table for the device given as a parameter. Note that it is not possible to use a parameter such as `/dev/nand0a` because this would instruct the command to access the RAM-partition table for the nand0a partition, but partitions within partitions are not supported.

```
    ptype  pname    pstart      plength
a:   ff      a    0xffffffff  0xffffffff
b:   6c      b    0x00000001  0x00000800
c:   ff      c    0xffffffff  0xffffffff
d:   ff      d    0xffffffff  0xffffffff
```

Note that the "ptype" output for the nand0b is 0x6c, this verifies that the second entry is filled. However, this has no further significance for the user; it is only used to indicate what partitions have been created with the 'part-add' command.

Now that the partition has been created, it can be accessed. Returning to the same example in the previous section, executing the following will now succeed:

```
losh> hd /dev/nand0b 512
```

You can create up to four partitions on a device; these partitions cannot overlap.

10.3 Partition Removal from RAM-Partition Table

Partitions can be removed with the 'part-rem' command. For example:

```
losh> part-rem /dev/nand0 b
```

will remove the second partition entry from the RAM-partition table. Removing this partition means that the contents will be overwritten with 0s.

For example:

```
info part /dev/nand0
```

will output the following table:

```
    ptype  pname    pstart      plength
a:   ff      a    0xffffffff  0xffffffff
```



```
b:   0       b    0x00000000  0x00000000
c:   ff      c    0xffffffff  0xffffffff
d:   ff      d    0xffffffff  0xffffffff
```

This table shows that the second partition entry has been removed (it is all 0s). Attempting to access this device will return an error message that the partition does not exist.

10.4 Writing RAM-Partition Table on the Device

Up to this point, we have only discussed the RAM-partition table that is created at boot. This table can be written from RAM to the device by using the 'part-write' command. The advantage of the 'part-write' command is that the partition table will reside on the device and, at boot, LogicLoader will fill up the RAM-partition table based on what is on the device. This allows you to save your partitions and not have to create them every time after a reset. The downside of writing the partition table to the device is that the table will take up some space in the first block; therefore, the entirety of the first block will not be available.

To write the RAM-partition table to the device, first create a partition. For example

```
losh> part-add /dev/nand0 b 1 2048
```

then type:

```
losh> part-write /dev/nand0
```

This command will write the RAM-partition table on the device (if the device-partition table is not empty, then the part-write command will fail). You can verify this by printing out the end of the 0th block. For example (on a NAND device):

```
hd /dev/nand0 512 16368
```

You will see the partition table at the end of the 512 bytes; note the 0x55aa partition marker at the end of the device. If at boot, the 0x55aa marker is found, then the device-partition table will be copied to RAM. If the partition marker is not found, then the RAM-partition table will be created at boot as normal.

Before writing the RAM-partition table to a NOR or NAND device, the portion of the block that will contain the partition table should only contain 1s (i.e., that portion of the block is in its default state, having never been written to; or that portion of the block has been erased, resetting all bits to 1). An exception to this rule is if additional information has been added to the end of the RAM-partition table and you want to write this additional information to the device-partition table. Provided that everything up to the point of that additional information is the same on the two partition tables, you can write to the device without first erasing the block. If a conflict between the two tables does exist, you will receive an error message when trying to perform the part-write command.

11 File Systems

11.1 File System Types

Two file system types are supported in LogicLoader: fat and YAFFS. YAFFS can only be mounted on NOR and NAND flash devices, while fat file systems (fatfs) can only be mounted on ATA devices (e.g., CompactFlash cards) and SD cards.

11.2 Mount Command

The general syntax of the 'mount' command is:

Mount <filesystem type> <device> <mount point> <flags>

File systems can only be mounted on partitions. An exception to this is YAFFS, which can be mounted on the entire device; although this usage is discouraged. However, if this scenario is required, a partition should be created that covers the entire device and then YAFFS should be mounted on this partition.

Specific examples of the 'mount' command will be presented in the sections below.

11.3 Mounting fatfs

An example of mounting a fat file system on an ATA device is:

```
losh> mount fatfs dev/ata0a /cf
```

This command will create a fat file system on the first partition of the ATA device.

If you would like to create a fat file system on an SD card:

```
losh> mount fatfs dev/sdmmc0a /sd
```

This command will create a fat file system on the first partition of the SD card.

In the examples above, the fat file system will be read only. If you would like to create a read/write file system you have to add a '–rw' flag at the end of the command line. For example:

```
losh> mount fatfs dev/ata0a /cf -rw
```

will create a read/write fat file system on the first partition of the ATA device.

11.3.1 Legacy Syntax

The old syntax of the 'mount' command is still supported, but it has many limitations. Because of these limitations, it is strongly recommended that you use the new syntax whenever possible. For example, the following command will create a fat file system on the first partition of the ATA device, just like the example above:

```
losh> mount fatfs /cf
```

However, the old syntax does not allow you to mount on any other partition other than the first. Also, the old syntax is only available for ATA devices; it cannot be used to mount file systems on SD cards.

11.4 Mounting YAFFS

11.4.1 Mounting YAFFS on NAND

An example of mounting a YAFFS file system on a NAND flash is:


```
losh> mount yaffs dev/nand0a /yaffs1
```

This command will create a YAFFS file system on the first partition of the NAND device. Before executing the 'mount' command, the partition should be created first. So the entire sequence would look like this:

```
losh> part-add dev/nand0 a 1 2048
losh> erase /dev/nand0a B0 B2048
losh> mount yaffs dev/nand0a /yaffs1
```

11.4.2 Mounting YAFFS on NOR

In order to mount YAFFS on a NOR flash device, an emulation layer must be created first. NOR flash devices do not have chunks, so a layer is required that emulates the NAND storage type making the NOR device look like a NAND device. The emulation layer is created using the 'mount' command.

```
losh> mount emu dev/flash0a /femu
```

Just like with NAND, mounting YAFFS on NOR requires a partition to be created first. So the entire sequence of commands to mount YAFFS on a NOR device looks like:

```
losh> part-add dev/flash0 a 5 20
losh> erase /dev/flash0a B0 B20
losh> mount emu dev/flash0a /femu1
losh> mount yaffs /femu1 /yaffs1
```

11.4.3 Legacy Syntax

The legacy 'add-yaffs' command is only supported for compatibility reasons and should not be actively used. The legacy command is emulated, so when you type `add-yaffs boot nand 1 2048` it creates a partition, registers the name "boot" for the actual partition. This means when you type `mount yaffs /boot` it gets the partition name based on '/boot' and performs a mount just like `mount yaffs /dev/nand0b /boot`.

12 YAFFS (Yet Another Flash File System)

Please be aware that the YAFFS user interface has changed in LogicLoader version 2.4. The legacy user interface is still supported for backwards compatibility, but the new interface—as described in this section—should be used whenever possible.

12.1 YAFFS Overview

The acronym YAFFS stands for the phrase "Yet Another Flash Filing System." YAFFS was developed by a company named Aleph One Limited and incorporated by Logic Product Development into the LogicLoader software program.

Logic selected YAFFS to fill its file system requirements due to the flexible nature of the program, its licensing scheme, and the fact that it is available for Linux, Windows CE, and other operating systems. YAFFS also allows LogicLoader and an RTOS to view and modify the same partition. It also makes it easier for customers to work with embedded flash technology and perform in-field updates. As an example, in Linux it is customary to have the Linux kernel reside in /boot/vmlinux, so using the commands below allows LogicLoader to mount, load, and boot the Linux kernel from the partition that is accessible from the Linux kernel.

```
losh> part-add /dev/nand0 a 9 500
losh> mount yaffs /dev/nand0a /nand-root
losh> load elf /nand-root/boot/vmlinux
losh> exec
```

Note: The partition entries for YAFFS partitions are not persistent—*they must be restored on each boot*. However, the partitions and data remain persistent.

12.2 Working with YAFFS in LogicLoader

12.2.1 Developing a Partition Scheme

In LogicLoader, YAFFS is mounted on partitions; there can be up to four partitions at a time on a NAND or NOR device.

Partitions are created with the 'part-add' command, as in the examples below. (Partition handling is discussed in detail in Section 10 of this document.) Customers should design a partitioning scheme which suits their individual needs; however, for the purpose of providing examples within this document, the following partitioning scheme will be assumed for NOR flash:

- A partition named "boot" which contains a bitmap and operating system image and spans the address space below:
  - * start: block 5
  - * length: 5 blocks (320kB)
- A partition named "data" which contains customer specific data.
  - * start: block 10
  - * length: 16 blocks (1 MB)

These two partitions are created with the following commands:

```
losh> part-add /dev/flash0 a 5 5
losh> part-add /dev/flash0 b 10 16
```

For the purpose of providing examples within this document, the following partitioning scheme will be assumed for NAND flash:

- A partition named "boot" which contains a bitmap and operating system image and spans the block range below:
  - * start: block 10

- * length: 256 blocks (8 MB, assuming 16kB block size)
- A partition named "data" which contains customer specific data.
  - * start: block 266 (abuts boot partition)
  - * length: 128 blocks (4 MB, assuming 16kB block size)

These two partitions are created with the following commands:

```
losh> part-add /dev/nand0 a 10 256
losh> part-add /dev/nand0 b 266 128
```

12.2.2 Formatting YAFFS Partitions

All file systems need to be formatted before they can be mounted. Because YAFFS was designed from the ground up to work with embedded flash technologies, it understands an 'erased' flash device to be both formatted and empty. To prepare your partition for mounting, use LogicLoader's 'erase' command to erase the area of flash where the partition is to be located.

Using the example partition scheme in the "Developing a Partition Scheme" section above, the partitions could be prepared for initial use by erasing the regions of the flash device spanned by them.

For a NOR example

```
losh> erase /dev/flash0 B5 B5
losh> erase /dev/flash0 B10 B16
```

For a NAND example:

```
losh> erase /dev/nand0 B10 B256
losh> erase /dev/nand0 B266 B128
```

Warning: Erasing flash blocks that will be used for YAFFS partitions will erase everything in those areas of flash. It is not required to format the partition every time the device is rebooted. The partition should only be formatted when an entirely new YAFFS partition is created, or when the data on a stored partition needs to be completely erased. For NAND-based devices, the first few blocks of NAND (the actual number of blocks is dependent on the NAND device) are used to hold the '/lboot' partition which is where LogicLoader resides. Modifying data in this partition can cause the board to fail to boot.

12.2.3 Mounting the Partition

To mount a partition, the 'mount' command is used—as was discussed in Section 11.4. This command takes the following arguments:

- <filesystem type> the type of file system being mounted ('yaffs' here)
- <device> the device on which the YAFFS partition is mounted
- <mountpoint> the name of the YAFFS partition

For example, to mount YAFFS on a NAND:

```
losh> mount yaffs /dev/nand0a /boot
losh> mount yaffs /dev/nand0b /data
```

Note: As previously discussed, you cannot mount YAFFS directly on NOR flash devices. First you have to mount an emulation layer on top of the NOR flash device, then you can mount YAFFS on the emulation layer.

For example:

```
losh> mount emu /dev/flash0a /femu1
```



```
losh> mount emu /dev/flash0b /femu2

losh> mount yaffs /femu1 /boot
losh> mount yaffs /femu2 /data
```

12.2.4 Accessing YAFFS Partitions in an OS

A key advantage of the read/write YAFFS file system capability at the LogicLoader level is the ability to share data stored in the file system with an OS environment. If an OS environment (e.g., Linux, Windows CE, VxWorks) implements YAFFS as an OS-accessible file-system, any files available to LogicLoader are also available to the OS, and vice-versa.

This contributes to significant benefits in the areas of system software upgrades (including OS upgrades) splash screen changes, script modifications, and other boot-time data that may need to be updated.

12.3 Summary

To use the YAFFS file system within LogicLoader, follow these steps:

1. Decide on a partitioning scheme and create partitions.
2. Format the partitions by erasing the associated flash blocks.
3. Mount the partitions using the 'mount' command (first mount an emulation layer for NOR flash devices).

Steps 1 and 3 must be repeated every time the system is booted. If the YAFFS partitions are frequently accessed, consider implementing steps 1 and 3 via a boot script. Step 2 only needs to be performed when creating a brand new partition or when the contents of an existing partition need to be completely erased.

Note: A partition is persistent. Re-adding a partition at boot-time restores access to previously saved data. Flash blocks must be erased to permanently remove a partition; otherwise, it can be recovered across boots.

Keep in mind the following when working with YAFFS and LogicLoader:

- Ensure partitions do not overlap each other, LogicLoader, or the configuration block.
- Ensure that a partition is erased before it is mounted for the first time.

Note: The legacy YAFFS mounting scheme is still supported for backwards compatibility, but its use is discouraged.

Appendix: LwIP License Agreement

LogicLoader uses the open source LwIP stack for networking support. The LwIP license requires the inclusion of the following license to satisfy Condition #2 below:

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This file is part of the lwIP TCP/IP stack.

Author: Adam Dunkels <adam@sics.se>