

Chapter 8: System Debugging and Profiling

8.1 Overview

In embedded systems and Linux-based environments, debugging and profiling are essential practices to ensure the proper functioning and optimal performance of the system. Debugging helps identify and resolve issues like incorrect behavior, crashes, or unexpected outputs, while profiling helps monitor the performance of the system, optimizing for resource usage such as CPU, memory, and I/O.

This chapter covers various techniques, tools, and methodologies for **debugging** and **profiling** Linux-based systems, particularly in the context of embedded systems.

8.2 System Debugging in Linux

System debugging involves identifying and fixing issues in both kernel and user-space components of the system. Linux provides a wide range of tools and techniques to troubleshoot problems, ranging from simple logs to advanced kernel debugging methods.

8.2.1 Debugging User-Space Applications

1. GDB (GNU Debugger):

- **GDB** is a powerful debugger for Linux systems that allows you to inspect and control the execution of a program. It can help diagnose issues like segmentation faults, memory access errors, and logical bugs.

2. Key Features:

- Set breakpoints to pause the execution at a specific line of code.
- Step through code line-by-line to inspect the flow.
- Examine and modify variables in memory during execution.

Basic Usage:

```
gdb ./my_program
```

```
(gdb) break main
```

```
(gdb) run
```

(gdb) backtrace # Get a stack trace of function calls

(gdb) print my_variable # Inspect a variable's value

(gdb) quit # Exit the debugger

3.

4. **Valgrind:**

- **Valgrind** is a memory debugging tool for detecting memory leaks, memory corruption, and memory management errors in C/C++ applications.

5. **Key Features:**

- Detects invalid memory access (e.g., accessing freed memory).
- Identifies memory leaks by tracking unfreed memory allocations.
- Provides a detailed log of memory usage to help pinpoint errors.

Basic Usage:

valgrind --leak-check=full ./my_program

6.

7. **Strace:**

- **Strace** is used to trace system calls made by a user-space process, providing insights into the interactions between the process and the kernel.

8. **Key Features:**

- Trace system calls such as file operations (**open**, **read**, **write**), network interactions, process creation, and more.
- Helps identify which system resources are being accessed and how.

Basic Usage:

strace ./my_program # Trace all system calls made by the program

9.

10. Core Dumps:

- A **core dump** is a file that captures the memory of a process at the time of a crash. Analyzing core dumps can help diagnose issues in applications that unexpectedly terminate.

11. Basic Usage:

Enable core dumps with:

```
ulimit -c unlimited # Allow core dumps
```

-
- Once the program crashes, a core dump file will be generated (e.g., **core**).

Analyze the core dump with GDB:

```
gdb ./my_program core
```

```
(gdb) backtrace # Get the stack trace
```

-

8.2.2 Debugging Kernel-Space Code

1. dmesg (Diagnostic Messages):

- The **dmesg** command displays kernel log messages, which provide information about system events, hardware interactions, driver messages, and errors.

2. Key Features:

- View real-time kernel logs, especially useful for debugging hardware or driver issues.
- Inspect the output from device drivers, kernel modules, and system processes.

Basic Usage:

```
dmesg | tail # View the most recent kernel messages
```

3.

4. Kernel Debugging with printk:

- The `printk` function is the kernel's equivalent of `printf`. It outputs messages to the kernel log (`dmesg`).

Usage Example:

```
printk(KERN_INFO "This is a debug message: %d\n", my_variable);
```

5.

6. KGDB (Kernel GNU Debugger):

- **KGDB** is a debugger for the Linux kernel. It allows you to debug kernel code on a live system, either via a serial connection or over a network.

7. Key Features:

- Set breakpoints in the kernel code.
- Examine kernel variables and memory.
- Step through kernel code to debug drivers or subsystems.

8. **Usage:** Requires setting up a remote debugging environment, typically over a serial cable or TCP/IP.

9. KDB (Kernel Debugger):

- **KDB** is a basic kernel debugger that can be used in environments where KGDB might not be available or suitable. It allows debugging via the console.

10. Usage:

- Activate KDB from the kernel command line by appending `kdb` to the boot parameters.
- Once activated, KDB provides commands for inspecting and modifying kernel memory, stack traces, and more.

8.3 System Profiling in Linux

Profiling helps you analyze the performance of a system or application by measuring metrics such as CPU usage, memory consumption, and I/O operations. It allows you to identify performance bottlenecks, optimize resource usage, and improve system efficiency.

8.3.1 Profiling User-Space Applications

1. **gprof:**

- **gprof** is a profiling tool that generates execution statistics, helping developers identify performance bottlenecks in user-space applications.

2. **Key Features:**

- Measures how much time is spent in each function.
- Helps identify which functions consume the most resources.

3. **Basic Usage:**

Compile with profiling enabled:

```
gcc -pg -o my_program my_program.c
```

○

Run the program to generate profiling data:

```
./my_program
```

○

View the profiling results:

```
gprof my_program gmon.out > analysis.txt
```

○

4. **perf:**

- **perf** is a performance analysis tool that provides detailed information about CPU performance counters, cache misses, context switches, and more.

5. **Key Features:**

- Collects hardware-level performance data like CPU cycles and cache misses.

- Supports both user-space and kernel-space profiling.

Basic Usage:

`perf stat ./my_program` # Get basic performance statistics

`perf record ./my_program` # Collect detailed profiling data

`perf report` # Generate a human-readable report

6.

7. **valgrind --tool=callgrind:**

- **Callgrind** is a profiling tool in Valgrind that focuses on the call graph, helping you visualize the performance of different function calls in your application.

Usage:

`valgrind --tool=callgrind ./my_program`

8. After execution, you can use **kcachegrind** to visualize the profiling data.

8.3.2 Profiling Kernel-Space Code

1. **ftrace:**

- **ftrace** is a powerful tracing utility built into the Linux kernel. It is used to trace kernel function calls, enabling performance analysis and debugging.

2. **Key Features:**

- Trace function calls in kernel space.
- Measure function execution times and track the execution flow.

3. **Basic Usage:**

Enable function tracing:

`echo function > /sys/kernel/debug/tracing/current_tracer`

```
echo 1 > /sys/kernel/debug/tracing/tracing_on
```

-

View trace output:

```
cat /sys/kernel/debug/tracing/trace
```

-

4. **perf for Kernel Profiling:**

- The **perf** tool also supports kernel profiling, providing insights into kernel execution at the function level, along with performance data such as cache misses and context switches.

Basic Usage:

```
perf record -e cpu-clock -a
```

```
perf report # View the profiling results
```

5.

6. **SystemTap:**

- **SystemTap** is another powerful tool that allows you to trace kernel functions and monitor kernel events dynamically.

7. **Key Features:**

- Trace kernel function calls and variables.
- Collect detailed data about kernel events and performance.

Basic Usage:

```
stap -v my_trace_script.stp
```

8.

8.4 Conclusion

Debugging and profiling are essential practices in Linux-based embedded systems development. By using the appropriate debugging tools like **GDB**, **strace**, and **dmesg**, you can identify and resolve issues in both user-space and kernel-space code. Profiling tools such as **perf**, **gprof**, and **valgrind** help optimize performance, improve resource usage, and ensure that the system is operating efficiently.

Mastering these tools and techniques will significantly improve your ability to develop, troubleshoot, and optimize embedded systems running Linux.