



# Outline

- 1 Foundations: Core Concepts
- 2 Theory: Evolution of Linux Threading
- 3 Mechanism: The `clone()` System Call
- 4 Kernel API & Lifecycle Management
- 5 Project Topic 2 Requirements

# Outline

- 1 Foundations: Core Concepts
- 2 Theory: Evolution of Linux Threading
- 3 Mechanism: The `clone()` System Call
- 4 Kernel API & Lifecycle Management
- 5 Project Topic 2 Requirements


# Foundations: Process vs. Thread

## Process (Resource Container)

- An instance of a running program.
- **Isolated Resources:** Owns Memory (Page Tables), File Descriptors, Signals.
- **Heavyweight:** Creation requires duplicating the entire address space (COW helps, but page tables are still copied).

## Thread (Execution Unit)

- "Lightweight Process" (LWP).
- **Shared Resources:** Code, Global Data, Heap, Open Files.
- **Private Resources:** Program Counter (PC), Register Set, Stack.



# Foundations: User-Level Threads (ULT)

**Definition:** Threads managed entirely by a user-space library (e.g., Green Threads, GNU Pth). The Kernel knows nothing about them.

## Mechanism

- Thread Table is stored in process memory.
- **Context Switch:** Just saving registers to user stack. No Mode Switch (User  $\leftrightarrow$  Kernel).
- **Speed:** Extremely fast (10 $\times$  faster than KLT).

## The Fatal Flaw

**Blocking System Calls:** If one ULT executes a blocking syscall (e.g., `read()`), the **Kernel blocks the entire process**.

*Result: All other threads stop, even if they are ready to run.*

# Foundations: Kernel-Level Threads (KLT)

**Definition:** Threads managed directly by the OS Kernel. (This is what we use in modern Linux).

## Mechanism

- Kernel maintains a Thread Table (TCB) for every thread.
- **Scheduling:** Kernel schedules threads, not processes.
- **Parallelism:** Can run on different CPU cores simultaneously.

## Trade-offs

- **Pros:** If one thread blocks, others continue running. True Multi-core utilization.
- **Cons:** Context switch requires a **Mode Switch** (User → Kernel → User), which is expensive (cache pollution, TLB implications).

# Outline

- 1 Foundations: Core Concepts
- 2 Theory: Evolution of Linux Threading
- 3 Mechanism: The `clone()` System Call
- 4 Kernel API & Lifecycle Management
- 5 Project Topic 2 Requirements

# The Big Picture: Threading Models

Before diving into Linux, we must understand the three mapping models:

## 1. Many-to-One (M:1)

- User-Level Threads.
- Kernel sees 1 process.
- *Example: Old Java Green Threads.*
- **Block one = Block all.**

## 2. One-to-One (1:1)

- **Linux Model (NPTL).**
- 1 User Thread = 1 Kernel Entity.
- *Example: Pthreads on Linux.*
- **True Parallelism.**

## 3. Many-to-Many (M:N)

- Hybrid approach.
- Complex scheduler in user space.
- *Example: Go Goroutines, Erlang.*
- **High concurrency, low overhead.**

# History: From "Hack" to Standard

## The Old Days: LinuxThreads (Pre-2.6)

- Used `clone()` but logic was flawed.
- Signal Handling Issue: Signals were sent to specific threads, not the process.
- PID Issue: Each thread had a different PID (broke POSIX compliance).

## The Modern Standard: NPTL (Native POSIX Thread Library)

- Introduced in kernel 2.6.
- 1:1 Model: Kernel manages scheduling directly.
- Solved the PID/Signal issues using Thread Group ID (Tgid).
- Heavy optimization for context switching.

# The Core Concept: PID vs. Tgid

Crucial for your understanding of kernel internals:

- **User Space View (getpid()):**
  - All threads in a process share the **Same PID**.
- **Kernel View (task\_struct):**
  - Each thread is a "Task".
  - Each task has a unique ID: pid (userspace calls this TID).
  - All threads share a tgid (Thread Group ID).

**Logic:** If pid == tgid, this task is the Main Thread.

```
1 struct task_struct {  
2     pid_t pid; // Unique per thread  
3     pid_t tgid; // Shared by group  
4  
5     struct task_struct *group_leader  
6     ;  
7     struct list_head thread_group;  
8 };
```

# Outline

- 1 Foundations: Core Concepts
- 2 Theory: Evolution of Linux Threading
- 3 Mechanism: The `clone()` System Call
- 4 Kernel API & Lifecycle Management
- 5 Project Topic 2 Requirements

# Creating Processes vs. Threads

In Linux, there is no "Create Thread" system call. There is only "Create Task".

System Call	Semantics	Sharing
fork()	Copy-on-Write (COW)	Shares nothing (mostly). Copies page tables.
vfork()	Block Parent	Shares memory (legacy, dangerous).
clone()	<b>Flexible</b>	<b>Selective sharing</b> via flags.

**Project Insight:** When you implement "Kernel Threads", you are essentially creating a task that shares **Kernel Memory** but has its own stack.

# Deep Dive: clone() Flags

Understanding these flags is required to understand kernel/fork.c.

```
1 // Essential Flags for Threads (pthread_create uses these)
2 #define CLONE_VM      0x00000100 // Share Memory Descriptor (mm_struct)
3 #define CLONE_FS      0x00000200 // Share Filesystem info (cwd, root)
4 #define CLONE_FILES    0x00000400 // Share File Descriptor Table (fd)
5 #define CLONE_SIGHAND  0x00000800 // Share Signal Handlers
6 #define CLONE_THREAD   0x00010000 // Put in same thread group (Same TGID)
7
```

## Question for Class

If I call `clone(CLONE_VM | CLONE_FILES)` but **NOT** `CLONE_THREAD`, what do I get?

# Deep Dive: clone() Flags

Understanding these flags is required to understand kernel/fork.c.

```
1 // Essential Flags for Threads (pthread_create uses these)
2 #define CLONE_VM      0x00000100 // Share Memory Descriptor (mm_struct)
3 #define CLONE_FS      0x00000200 // Share Filesystem info (cwd, root)
4 #define CLONE_FILES    0x00000400 // Share File Descriptor Table (fd)
5 #define CLONE_SIGHAND  0x00000800 // Share Signal Handlers
6 #define CLONE_THREAD   0x00010000 // Put in same thread group (Same TGID)
7
```

## Question for Class

If I call `clone(CLONE_VM | CLONE_FILES)` but **NOT** `CLONE_THREAD`, what do I get? **Answer: A classic "Lightweight Process" (LWP) that shares memory but has a different PID in userspace (like old LinuxThreads).**

# Outline

- 1 Foundations: Core Concepts
- 2 Theory: Evolution of Linux Threading
- 3 Mechanism: The `clone()` System Call
- 4 Kernel API & Lifecycle Management
- 5 Project Topic 2 Requirements

# API: Creating Kernel Threads

Unlike user threads, Kernel Threads (kthreads) have no User Space memory (mm is NULL).

**Header:** <linux/kthread.h>

## 1. The Easy Way: kthread\_run

```
1 struct task_struct *ts;
2 // Creates thread AND wakes it up immediately
3 ts = kthread_run(thread_fn, data, "worker_%d", id);
4
```

## 2. The Manual Way: kthread\_create

```
1 ts = kthread_create(thread_fn, data, "worker");
2 if (!IS_ERR(ts)) {
3     // You can bind CPU affinity here before starting!
4     kthread_bind(ts, cpu_id);
5     wake_up_process(ts);
6 }
7
```

# API: Stopping Kernel Threads

**Warning:** Killing kernel threads forcefully is bad. They must exit voluntarily.

## Worker Logic (Thread Function):

```
1 int thread_fn(void *data) {
2     while (!kthread_should_stop()) {
3         // Do work...
4
5         // IMPORTANT: Yield CPU
6         schedule_timeout_interruptible(HZ);
7     }
8     return 0;
9 }
10
```

## Controller Logic (Module Exit):

```
1 // Request the thread to stop
2 int ret = kthread_stop(ts);
3 // This function blocks until thread exits!
4
```

# API: CPU Affinity (Project Requirement)

One of your tasks is to monitor/set CPU Affinity.

## Setting Affinity (Binding)

```
1 // Bind current thread to CPU 0
2 kthread_bind(current, 0);
3
```

## Checking Affinity (Statistics)

Inside task\_struct, look for:

```
1 struct cpumask cpus_mask; // Allowed CPUs
2 int cpu; // Current CPU
3
```

*Note: You might need to export this via /proc or printk to verify your scheduler is working.*

# Outline

- 1 Foundations: Core Concepts
- 2 Theory: Evolution of Linux Threading
- 3 Mechanism: The `clone()` System Call
- 4 Kernel API & Lifecycle Management
- 5 Project Topic 2 Requirements

# Basic Requirements: Implementation Guide

**Goal:** Create a Kernel Module that spawns threads and measures them.

## ① Module Init:

- Create  $N$  kernel threads using `kthread_create`.
- Bind them to specific Cores (e.g., Thread 0 → CPU 0).

## ② The Payload:

- The threads should do something measurable (e.g., increment a shared atomic counter, or just sleep/wake).

## ③ Measurement (The "Analysis" part):

- **Context Switches (CSW):** Read `current->nvcsw` (Voluntary) and `current->nivcsw` (Involuntary).
- Log these values to `dmesg` or a custom `/proc` file before exit.

# Requirement: KLT vs ULT Comparison

You must compare your Kernel Module threads against a User-Space Pthread program.

## Experiment Setup:

- **Scenario A:** 4 Threads on 4 Cores incrementing a shared atomic variable.
- **Scenario B:** High-frequency yielding (`yield()`).

## What to measure?

- **System Time (sys):** High for KLT (syscall overhead if communicating).
- **User Time (user):** Zero for KLT!
- **Throughput:** Operations per second.

# Advanced Option: User-Space Coroutines

**Challenge:** Implement "Green Threads" without Kernel support.

**Mechanism: Stack Switching**

- You need to allocate a memory block (heap) to act as a stack.
- Use `setjmp` / `longjmp` (C Library) OR `swapcontext` (`ucontext.h`).

```
1 #include <ucontext.h>
2 ucontext_t main_ctx, thread_ctx;
3 char stack[16384];
4
5 void thread_func() { ... }
6
7 // Setup
8 getcontext(&thread_ctx);
9 thread_ctx.uc_stack.ss_sp = stack;
10 thread_ctx.uc_stack.ss_size = sizeof(stack);
11 makecontext(&thread_ctx, thread_func, 0);
12
13 // Switch
14 swapcontext(&main_ctx, &thread_ctx);
15
```

# Advanced Option: Hybrid Interaction

**Goal:** Communicate between User Space and Kernel Thread.

**Ideas:**

- **Netlink Socket:** Asynchronous, message-based. Best for event notification.
- **Debugfs / Sysfs:** Use a file to pass commands.
- **Custom System Call:** (Hardcore) Add a syscall that wakes up a specific wait\_queue in the kernel.

**Use Case Example:** A user program captures network packets (raw socket) and passes them to a high-priority Kernel Thread for processing without copying data (Zero-Copy using shared memory).

# Evaluation & Grading (Topic 2 Specific)

---

Component	Detailed Focus
<b>Implementation (25%)</b>	Correct usage of <code>kthread_run/stop</code> . Accurate statistics (CSW, CPU Affinity). Clean cleanup (no zombies).
<b>Advanced Options (10%)</b>	Completeness of User-Space Coroutines (Stack switching) OR Hybrid Communication mechanism OR other innovation.
<b>Stability (5%)</b>	<b>Crucial:</b> No Kernel Panics during the live demo. System must remain responsive.
<b>Presentation (20%)</b>	<b>Demo (10%):</b> Successful live run. <b>Structure (5%):</b> Clarity of technical explanation. <b>Q&amp;A (5%):</b> Ability to answer kernel-level questions .
<b>Final Report (20%)</b>	<b>Guidelines (5%)</b> <b>Depth (10%)</b> <b>Analysis (5%)</b>

---

# Resources & Next Steps

## Recommended Reading:

- *Linux Kernel Development* (Robert Love) - Chapter 3 (Process Management).
- `man clone`, `man setjmp`.
- Source: `kernel/kthread.c`, `kernel/fork.c`.

## Immediate Action Items:

- ① Compile a "Hello World" kernel module.
- ② Add `kthread_run` to it.
- ③ Try to unload the module (`rmmod`). If it hangs, you forgot `kthread_should_stop()`!