

Principia Softwarica

Plan 9 Code Explained

Yoann Padioleau

IWP9



The Best Engineers Understand the Stack

Facebook: linker slow, but benchmark for everyone
→ most engineers: stack, accepted it, black box

We had to wait for [Ian Lance Taylor](#) (Google, 2006) to write `glibc` → a new linker for GNU Binutils.

Not a young hotshot. Old school Unix/GCC hacker (since 1996).
Also part of the Go team at Google.
The kind of engineer who understands the full stack.
Understand deeply = build better

The Education Gap

"What happens when you type `ls` in a terminal window?"

[Keith Adams](#) → colleague at Facebook, later Chief Architect at Slack → used this as his interview question.

Simple question. The answer involves: the shell, the C library, the kernel, the graphics stack, the windowing system.
Most engineers can't answer.

"Full stack" today = React + Node + cloud.
REAL full stack = compiler, linker, kernel, graphics

The AI Era Makes This More Important

We write less code. We need **MORE** understanding.

AI generates code → but who evaluates it?
Who debugs the segfault? The linker error?

AI tools like `CLAUDE` Code USE the classic commands:
`grep`, `ls`, `diff`, `awk`, `gcc`, `ls` ...
The AI is a power user of the very programs Principia explains.

Writing code = 20%. Understanding = 80%.
AI handles the 20%. The 80% becomes everything.

But the Stack is Enormous

| | |
|--------------|-----------|
| Linux kernel | ~40M LOC |
| glibc | ~15M LOC |
| bash | ~3.5M LOC |
| log server | ~100M LOC |
| stern alone | ~80M LOC |

Each component: more code than ALL of Plan 9

Plan 9 - 100k Lines Code

| | |
|-------------|----------|
| Kernel | ~10k LOC |
| glibc | ~10k LOC |
| bash | ~10k LOC |
| log server | ~10k LOC |
| stern alone | ~10k LOC |

Principia Software

A Series of Books

| Category | Book | Program(s) | LOC | LOE | Pages |
|-----------------------|------------------|---|---------------|--------------|-------------|
| Core system | Kernel | 9pi | 35000 | 3200 | 732 |
| | Core libraries | libc libregexp libthread | 19000 | 1600 | 438 |
| | Shell | rc | 6500 | 1700 | 166 |
| Development toolchain | C compiler | 5c libcc | 18500 | 1900 | 471 |
| | Assembler | 5a | 3600 | 4400 | 176 |
| | Linker | 5l | 7500 | 5400 | 296 |
| Developer tools | Editor | ed | 1600 | 200 | 45 |
| | Build system | mk | 4350 | 4050 | 197 |
| | Debuggers | db acid strace libmach | 13100 | 1000 | 321 |
| Graphics | Graphics stack | prof time kprof stats iostats | 3900 | 350 | 102 |
| | | /dev/draw | 18500 | 3400 | 507 |
| | | libdraw libmemdraw libmemlayer | | | |
| Networking | Windowing system | rio libframe libcomplete libplumb | 8800 | 4000 | 289 |
| | | /dev/net | 18300 | 4800 | 457 |
| | | libip lib9p | | | |
| Misc | CLI utilities | cat ls grep sed diff tar gzip ... | 23900 | 650 | 493 |
| Total: | | | 182550 | 36650 | 4690 |

Table 1: Principia Softwarica Books and Their Statistics.
(LOC = lines of code; LOE = lines of explanation; Pages = number of typeset pages)

A New Plan 9 Fork

```

-- assemblers/      -- kernel/          -- mkfiles/
-- 5a/              -- arch/            -- 386/
-- 8a/              -- concurrency/    -- arm/
-- data2s.c         -- console/         -- makedirs
-- Assembler.nw    -- devices/         -- mkfile.proto
-- mkfile           -- files/           -- mklib
-- builders/       -- filesystems/    -- mkone
-- Builder.nw       -- init/            -- mkfile-target-pc
-- mk/              -- memory/         -- mkfile-target-pi
-- mkfile           -- Kernel.nw       -- networking/
-- compilers/      -- mkfile           -- arp/
-- 5c/              -- network/         -- dhcp/
-- 8c/              -- processes/      -- ftp/
-- cc/              -- syscalls/       -- http/
-- Compiler.nw     -- lib_core/       -- ip/
-- cpp/            -- libbio/          -- mkfile
-- mkfile           -- libc/            -- ndb/
-- debuggers/     -- libcore.nw      -- Network.nw
-- acid/           -- libthread/      -- snoopy/
-- db/             -- mkfile           -- telnet/
-- Debugger.nw    -- lib_graphics/   -- profilers/
-- libmach/        -- Graphics.nw     -- iostats/
-- mkfile          -- libdraw/         -- mkfile
-- Dockerfile      -- libing/          -- Profiler.nw
-- docs/           -- libmemdraw/     -- shells/
-- articles/       -- libmenlayer/    -- mkfile
-- iw9p/           -- mkfile           -- rc/
-- dosdisk.img     -- lib_networking/ -- Shell.nw
-- editors/        -- lib9p/          -- utilities/
-- ed/             -- libip/           -- archive/
-- mkfile          -- mkfile           -- byte/
-- include/        -- lib_strings/    -- calc/
-- lib.h           -- libflate/        -- compare/
-- ...            -- libregexp/      -- files/
-- u.h            -- libstring/      -- mkfile
-- linkers/        -- mkfile           -- pipe/
-- 5l/             -- Makefile         -- process/
-- 8l/             -- mkconfig.pc     -- text/
-- Linker.nw       -- mkconfig.pi     -- time/
-- mkfile          -- mkfile           -- Utilities.nw
-- readme.txt     -- mkfile-host-Cygwin
-- windows/       -- mkfile-host-Linux
-- libcomplete/   -- mkfile-host-macOS
-- libplumb/      -- mkfile-host-Plan9
-- mkfile         -- MISC/
-- rio/           -- pc/
-- Windows.nw     -- pi/
  
```

Figure 1: <https://github.com/aryx/principia-softwarica> Layout.
(as generated by `tree -L 2 -F` and spread in 3 columns)

A New Way to Build

Goken Dockerfile excerpt

```

# Build goken on Ubuntu using
# gcc/binutils (and mk/rc)
FROM ubuntu:24.04

# Setup a basic C dev environment
RUN apt-get install -y gcc libc6-dev \
  byacc

# Now let's build from source
WORKDIR /src
COPY . .

# Small shell script (not GNU autoconf)
# to detect arch and generate
# ./mkconfig
RUN ./configure

# The script below obviously builds
# 'mk' (without needing mk) but also:
# - 'rc', which is called by 'mk'
# - 'ed', for the mkenam script
# run during the build
RUN ./scripts/build-mk.sh

# copy ./ROOT/<arch>/bin/{mk,rc,ed}
# to ./bin/
RUN ./scripts/promote-mk.sh

# make mk and rc accessible
ENV PATH="/src/bin:${PATH}"
ENV MKSHELL="/src/bin/rc"

# Let's build goken (using mk/rc built
# in the previous step)
RUN mk
RUN mk install

ENV PATH="/src/ROOT/amd64/bin:\
  /src/ROOT/arm64/bin:\
  ${PATH}"
  
```

Principia Softwarica Dockerfile

```

# Build principia on Ubuntu
# Linux (amd64 or arm64)
# for 386 (pc) and arm (pi)
FROM padator/goken

WORKDIR /principia
COPY . .

# 386
RUN cp mkconfig.pc mkconfig
RUN mk && mk install
RUN mk kernel

# arm
RUN cp mkconfig.pi mkconfig
RUN mk && mk install
RUN mk kernel

# VFAT tools
RUN apt-get install -y dosfstools \
  mtools

# making dosdisk.img
RUN dd if=/dev/zero of=tmp.img \
  bs=1M count=512
RUN mks.vfat tmp.img
RUN mcopy -i tmp.img -s -o \
  ROOT/* : :
RUN cat MISC/pc/bootsector tmp.img \
  > dosdisk.img
RUN rm -f tmp.img
  
```

Table 3: Principia Softwarica Build and Packaging Process.

The Best Engineers Understand the Stack

Facebook: linker slow, bottleneck for everyone
→ most engineers: stuck, accepted it, black box

We had to wait for **Ian Lance Taylor** (Google, 2008) to write `gold` — a new linker for GNU binutils.

Not a young hotshot. Old-school Unix/GCC hacker (since 1990).

Also part of the **Go team** at Google.

The kind of engineer who understands the full stack.

Understand deeply = build better

The Education Gap

"What happens when you type ls in a terminal window?"

Keith Adams — colleague at Facebook, later Chief Architect at Slack — used this as his interview question.

Simple question. The answer involves: the shell, the C library, the kernel, the graphics stack, the windowing system.

Most engineers can't answer.

“Full stack” today = React + Node + cloud

REAL full stack = compiler, linker, kernel, syscalls

The AI Era Makes This More Important

We write less code. We need MORE understanding.

AI generates code — but who evaluates it?

Who debugs the segfault? The linker error?

AI tools like Claude Code USE the classic commands:

`grep`, `ls`, `diff`, `awk`, `gcc`, `ld` ...

The AI is a power user of the very programs Principia explains.

Writing code = 20%. Understanding = **80%**.

AI handles the 20%. The 80% becomes everything.

But the Stack is Enormous

| | | |
|--------------|-----------|---------------------|
| Linux kernel | ~40M LOC | core ~2M |
| gcc | ~15M LOC | |
| glibc | ~1.5M LOC | |
| bash | ~180K LOC | |
| Xorg server | ~500K LOC | ecosystem: millions |
| xterm alone | ~88K LOC | |

Each component: more code than ALL of Plan 9



(Kernel)
Linux kernel
40M

(Compiler)

gcc
15M

(Libc)
glibc
1.5M

(Debugger)
gdb
1.5M

binutils
700K

Linux
net
stack
400K

vim
350K

1

2

(Graphics+Windowing)
X Window System
1.5M

Plan

3

4

5

1 bash 180K (Shell) 2 grep+sed+awk 120K 3 coreutils 100K 4 tar+diff 80K 5 make 35K

■ Plan 9 = 183K LOC – ALL of the above

Plan 9 — 100x Less Code

Thompson, Ritchie, Pike — Whole OS: ~183,000 LOC

Plan 9

rio = 8,800 LOC

terminal < 1K

/dev/cons

Linux / X

Xorg = millions

xterm = 88K

tty, pty, VT100...

Ideal teaching OS: Open source, Small, Real, Complete, Coherent, Self-hosting

Principia Softwarica

Research OS → Teaching OS

Books explaining **ALL** source code of **ALL** essential programs

The books **ARE** the implementations (literate programs)

kernel → compiler → assembler → linker

shell → libc → editor → build system

debugger → graphics → windowing → networking

"What happens when you type ls?"

With Principia: fully answerable.

Principia Softwarica

Plan 9 Code Explained

Yoann Padioleau

IWP9



The Best Engineers Understand the Stack

Facebook: linker slow, but benchmark for everyone
→ most engineers: stack, accepted it, black box

We had to wait for Ian Lance Taylor (Google, 2006) to write glibc → a new linker for GNU binutils.

Not a young hotshot. Old school Unix/GCC hacker (since 1996).
Also part of the Go team at Google.
The kind of engineer who understands the full stack.
Understand deeply = build better

The Education Gap

"What happens when you type ls in a terminal window?"

Kath Adcox → colleague at Facebook, later Chief Architect at Slack → used this as his interview question.

Simple question. The answer involves: the shell, the C library, the kernel, the graphics stack, the windowing system.
Most engineers can't answer.

"Full stack" today = React + Node + cloud.
REAL full stack = compiler, linker, kernel, graphics

The AI Era Makes This More Important

We write less code. We need MORE understanding.

AI generates code → but who evaluates it?
Who debugs the segfault? The linker error?

AI tools like CLAUDE: Code USE the classic commands: glibc, ls, diff, awk, gcc, fd ...
The AI is a power user of the very programs Principia explains.

Writing code = 20%. Understanding = 80%.
AI handles the 20%. The 80% becomes everything.

But the Stack is Enormous

| | |
|--------------|-----------|
| Linux kernel | ~40M LOC |
| glibc | ~15M LOC |
| bash | ~3.5M LOC |
| log server | ~500K LOC |
| stern alone | ~80K LOC |

Each component: more code than ALL of Plan 9

A Series of Books

| Category | Book | Program(s) | LOC | LOE | Pages |
|-----------------------|------------------|---|---------------|--------------|-------------|
| Core system | Kernel | 9pi | 35000 | 3200 | 732 |
| | Core libraries | libc libregexp libthread | 19000 | 1600 | 438 |
| | Shell | rc | 6500 | 1700 | 166 |
| Development toolchain | C compiler | 5c libcc | 18500 | 1900 | 471 |
| | Assembler | 5a | 3600 | 4400 | 176 |
| | Linker | 5l | 7500 | 5400 | 296 |
| Developer tools | Editor | ed | 1600 | 200 | 45 |
| | Build system | mk | 4350 | 4050 | 197 |
| | Debuggers | db acid strace libmach | 13100 | 1000 | 321 |
| | Profilers | prof time kprof stats iostats | 3900 | 350 | 102 |
| Graphics | Graphics stack | /dev/draw libdraw libmemdraw libmemlayer | 18500 | 3400 | 507 |
| | Windowing system | rio libframe libcomplete libplumb | 8800 | 4000 | 289 |
| Networking | Network stack | /dev/net libip lib9p | 18300 | 4800 | 457 |
| Misc | CLI utilities | cat ls grep sed diff tar gzip ... | 23900 | 650 | 493 |
| Total: | | | 182550 | 36650 | 4690 |

Table 1: Principia Softwarica Books and Their Statistics.
(LOC = lines of code; LOE = lines of explanation; Pages = number of typeset pages)

A New Plan 9 Fork

```

Why a New Fork?
https://github.com/aryx/principia-softwarica

Repository structure:
- 5a/
- 5l/
- data2s.c
- Assembler.nw
- mkfile
- Builder.nw
- mk/
- compilers/
- 5c/
- 8c/
- cc/
- Compiler.nw
- cpp/
- mkfile
- debuggers/
- acid/
- db/
- Debugger.nw
- libmach/
- mkfile
- Dockerfile
- docs/
- articles/
- iw9p/
- dosdisk.img
- editors/
- ed/
- mkfile
- include/
- lib.h
- ...
- u.h
- linkers/
- 5l/
- 5l/
- Linker.nw
- mkfile
- readme.txt
- windows/
- libcomplete/
- libplumb/
- mkfile
- rio/
- Windows.nw

kernel/
- arch/
- concurrency/
- console/
- devices/
- files/
- filesystems/
- init/
- memory/
- Kernel.nw
- mkfile
- network/
- processes/
- syscalls/
- lib_core/
- libbio/
- libc/
- Libcore.nw
- libthread/
- mkfile
- lib_graphics/
- Graphics.nw
- libdraw/
- libing/
- libmemdraw/
- libmenlayer/
- mkfile
- lib_networking/
- lib9p/
- libip/
- mkfile
- lib_strings/
- libflate/
- libregexp/
- libstring/
- mkfile
- Makefile
- mkconfig.pc
- mkconfig.pi
- mkfile
- mkfile-host-Cygwin
- mkfile-host-Linux
- mkfile-host-macOS
- mkfile-host-Plan9
- MISC/
- pc/
- pi/

mkfiles/
- 386/
- arm/
- makedirs
- mkfile.proto
- mklib
- mkone
- mkfile-target-pc
- mkfile-target-pi
- networking/
- arp/
- dhcp/
- ftp/
- http/
- ip/
- mkfile
- ndb/
- Network.nw
- snoopy/
- telnet/
- profilers/
- iostats/
- mkfile
- Profiler.nw
- shells/
- mkfile
- rc/
- Shell.nw
- utilities/
- archive/
- byte/
- calc/
- compare/
- files/
- mkfile
- pipe/
- process/
- text/
- time/
- Utilities.nw

ROOT/
- 386/
- arm/
- lib/
- rc/
- tests/
- usr/
    
```

Figure 1: <https://github.com/aryx/principia-softwarica> Layout.
(as generated by `tree -L 2 -F` and spread in 3 columns)

A New Way to Build

Goken Dockerfile excerpt

```

# Build goken on Ubuntu using
# gcc/binutils (and mk/rc)
FROM ubuntu:24.04

# Setup a basic C dev environment
RUN apt-get install -y gcc libc6-dev \
    byacc

# Now let's build from source
WORKDIR /src
COPY . .

# Small shell script (not GNU autoconf)
# to detect arch and generate
# ./mkconfig
RUN ./configure

# The script below obviously builds
# 'mk' (without needing mk) but also:
# - 'rc', which is called by 'mk'
# - 'ed', for the mkenam script
# run during the build
RUN ./scripts/build-mk.sh

# copy ./ROOT/<arch>/bin/{mk,rc,ed}
# to ./bin/
RUN ./scripts/promote-mk.sh

# make mk and rc accessible
ENV PATH="/src/bin:${PATH}"
ENV MKSHELL="/src/bin/rc"

# Let's build goken (using mk/rc built
# in the previous step)
RUN mk
RUN mk install

ENV PATH="/src/ROOT/amd64/bin:\
    /src/ROOT/arm64/bin:\
    ${PATH}"
        
```

Principia Softwarica Dockerfile

```

# Build principia on Ubuntu
# Linux (amd64 or arm64)
# for 386 (pc) and arm (pi)
FROM padator/goken

WORKDIR /principia
COPY . .

# 386
RUN cp mkconfig.pc mkconfig
RUN mk && mk install
RUN mk kernel

# arm
RUN cp mkconfig.pi mkconfig
RUN mk && mk install
RUN mk kernel

# VFAT tools
RUN apt-get install -y dosfstools \
    mtools

# making dosdisk.img
RUN dd if=/dev/zero of=tmp.img \
    bs=1M count=512
RUN mks.vfat tmp.img
RUN mcopy -i tmp.img -s -o \
    ROOT/* : :
RUN cat MISC/pc/bootsector tmp.img \
    > dosdisk.img
RUN rm -f tmp.img
        
```

Table 3: Principia Softwarica Build and Packaging Process.

A Series of Books

| Category | Book | Program(s) | LOC | LOE | Pages | |
|--------------------------|----------------|-------------|-------|------|-------|--|
| Core system | Kernel | 9pi | 35000 | 3200 | 732 | |
| | Core libraries | libc | 19000 | 1600 | 438 | |
| | | libregexp | | | | |
| | | libthread | | | | |
| | Shell | rc | 6500 | 1700 | 166 | |
| Development toolchain | C compiler | 5c libcc | 18500 | 1900 | 471 | |
| | Assembler | 5a | 3600 | 4400 | 176 | |
| | Linker | 5l | 7500 | 5400 | 296 | |
| Developer tools | Editor | ed | 1600 | 200 | 45 | |
| | Build system | mk | 4350 | 4050 | 197 | |
| | Debuggers | db acid | 13100 | 1000 | 321 | |
| | | strace | | | | |
| | | libmach | | | | |
| Profilers | prof time | 3900 | 350 | 102 | | |
| | | kprof stats | | | | |

| | | | | | |
|---------------|------------------|---|---------------|--------------|-------------|
| Graphics | Graphics stack | /dev/draw libdraw libmemdraw libmemlayer | 18500 | 3400 | 507 |
| | Windowing system | rio libframe libcomplete libplumb | 8800 | 4000 | 289 |
| Networking | Network stack | /dev/net libip lib9p | 18300 | 4800 | 457 |
| Misc | CLI utilities | cat ls grep sed diff tar gzip ... | 23900 | 650 | 493 |
| Total: | | | 182550 | 36650 | 4690 |

Table 1: Principia Softwarica Books and Their Statistics.
(**LOC** = lines of code; **LOE** = lines of explanation; **Pages** = number of typeset pages)

The Bookset

15 books, 7 categories, 183K LOC total
ARM architecture (simplest widely-used + Raspberry Pi)
LOE/LOC ~1 target — work in progress

The Utilities book: cat, ls, grep, sed, diff, tar ...
The exact commands AI coding tools run hundreds of times a day.



Literate Programming

Donald Knuth, 1984

Mix code + documentation
Present in order of UNDERSTANDING, not compiler order
Piece by piece, separate concerns

Tool: Noweb

<<chunk-name>> defines a chunk <<chunk-name>> references a chunk Chunks defined incrementally, across chapters

Literate programming 101



Literate Programming in Practice



Why Books?

"Are books still the best medium in 2026?"

Books: teaching for ~2000 years. Survived radio, TV, internet, MOOCs.

- IDE: compiler order, no narrative
- Videos: can't reference, can't search
- AI explanation: no coherent narrative

Syncweb brings IDE features INTO the book

Narrative of a book + navigation of an IDE



Transfer — Learn Here, Apply Everywhere

Understanding one small elegant OS gives deep intuition about Linux, macOS, and even Windows.

Plan 9 ideas already everywhere: UTF-8, /proc, Linux namespaces (=Docker), Go, SP in WSL2

Not just theory — you can use the tools:
plan9port: acme, ssh, rc, grep, sed, awk on Linux/macOS
gokernelcc: Plan 9 compilers that build Linux/macOS/Windows programs

Use them in your own environment, outside Plan 9.
The same tools AI coding agents run hundreds of times a day.

The Bookset

15 books, 7 categories, **183K LOC** total

ARM architecture (simplest widely-used + Raspberry Pi)

LOE/LOC ~1 target — work in progress

The **Utilities** book: `cat`, `ls`, `grep`, `sed`, `diff`, `tar` ...

The exact commands AI coding tools run hundreds of times a day.




How many 400-page books to explain it all?

each  = one 400-page book




Principia Softwarica covers ALL of Plan 9 in 15 books.

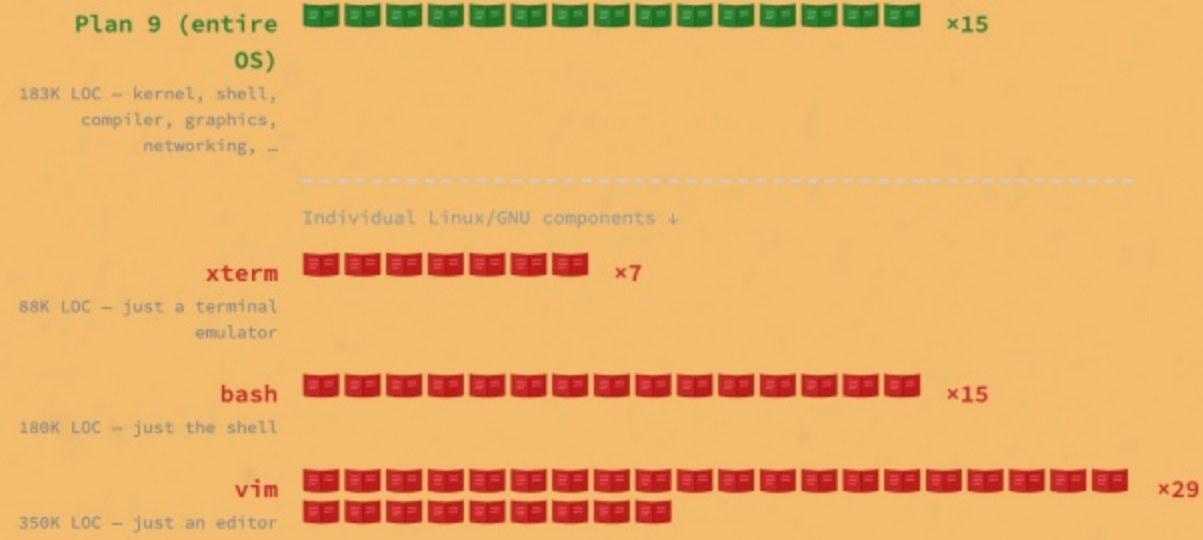
bash *alone* would need as many books as the entire Plan 9 OS.



same scale, zoomed out ↓


How many 400-page books to explain it all?

each  = one 400-page book

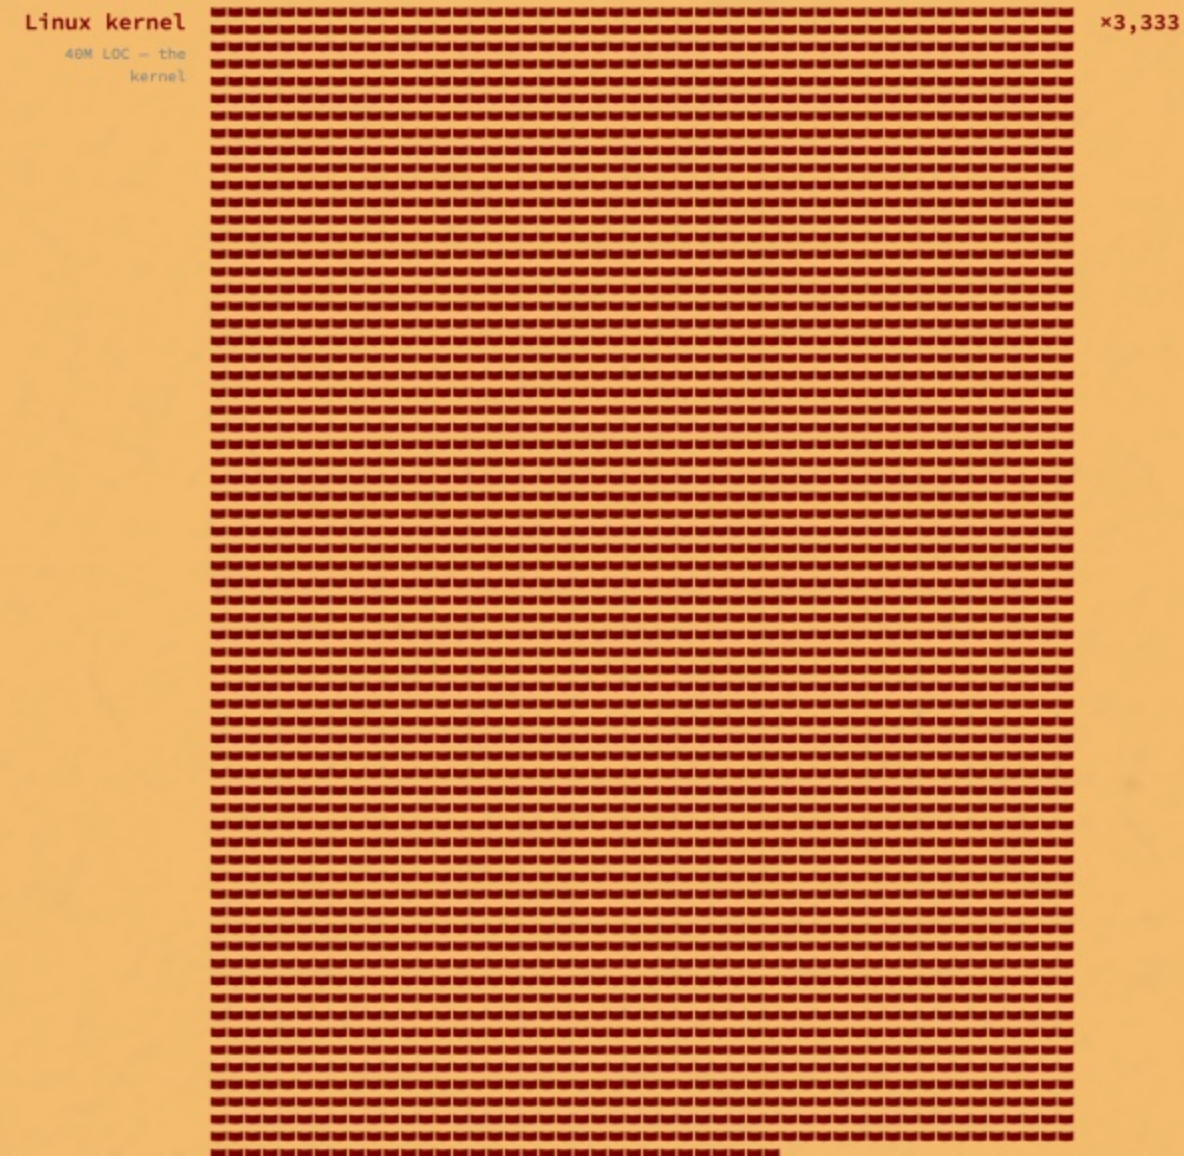
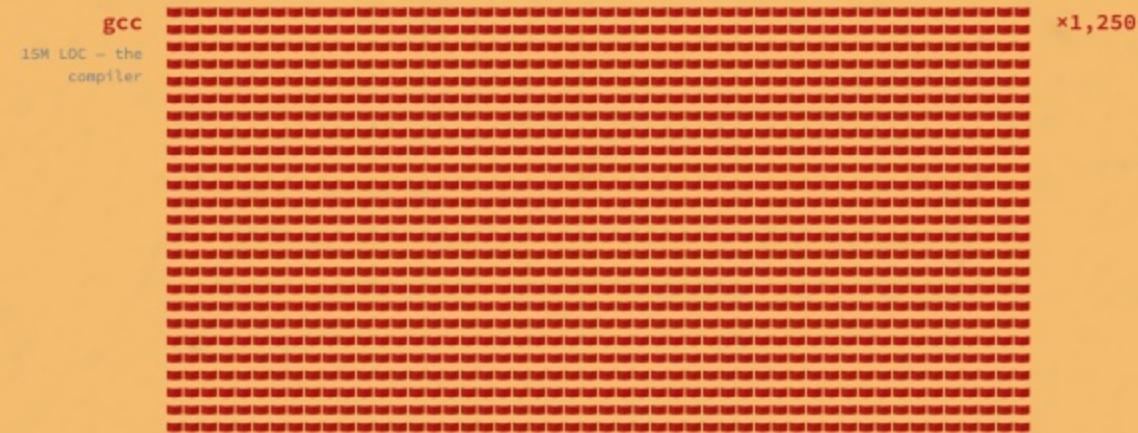


Principia Softwarica covers ALL of Plan 9 in 15 books.

bash *alone* would need as many books as the entire Plan 9 OS.



same scale, zoomed out ↓

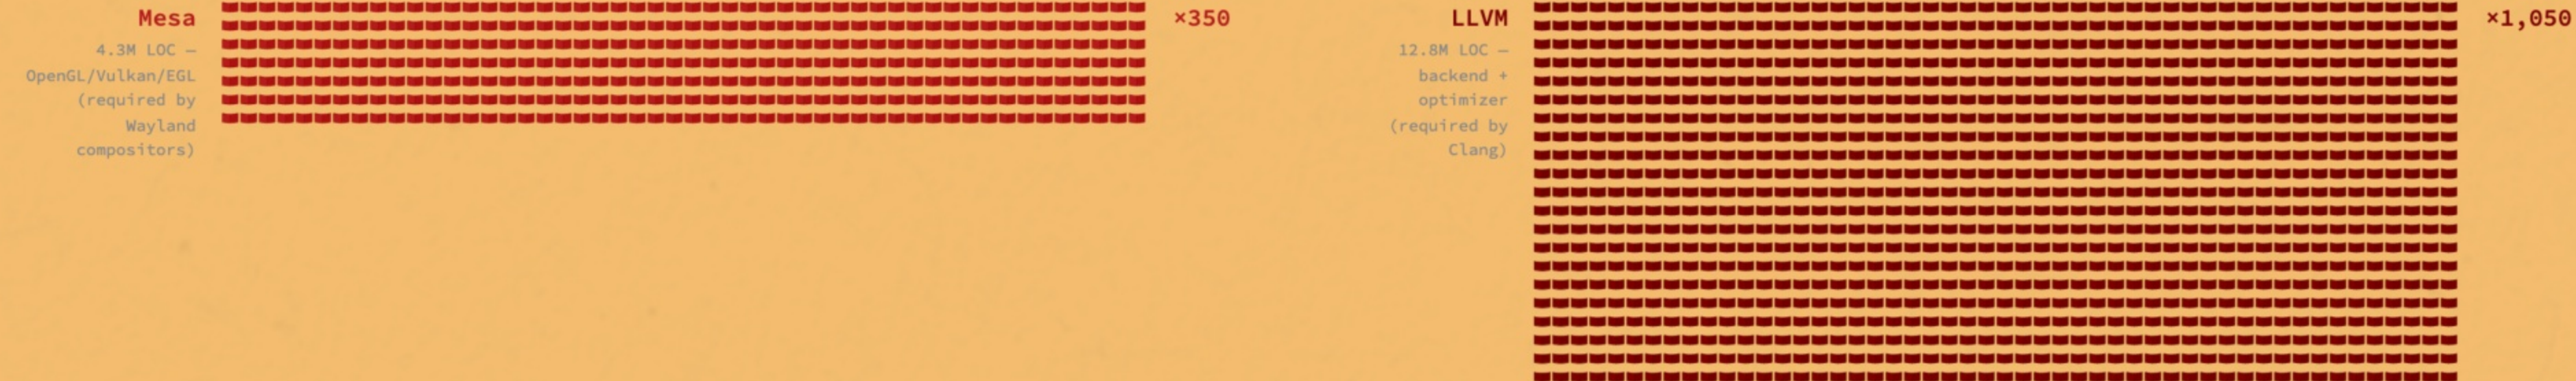


es a day.

"Clean rewrites" – Wayland replaces X11, Clang replaces gcc



... but they require these massive dependencies ↓



Literate Programming

Donald Knuth, 1984

Mix code + documentation

Present in order of **UNDERSTANDING**, not compiler order

Piece by piece, separate concerns

Tool: **Noweb**

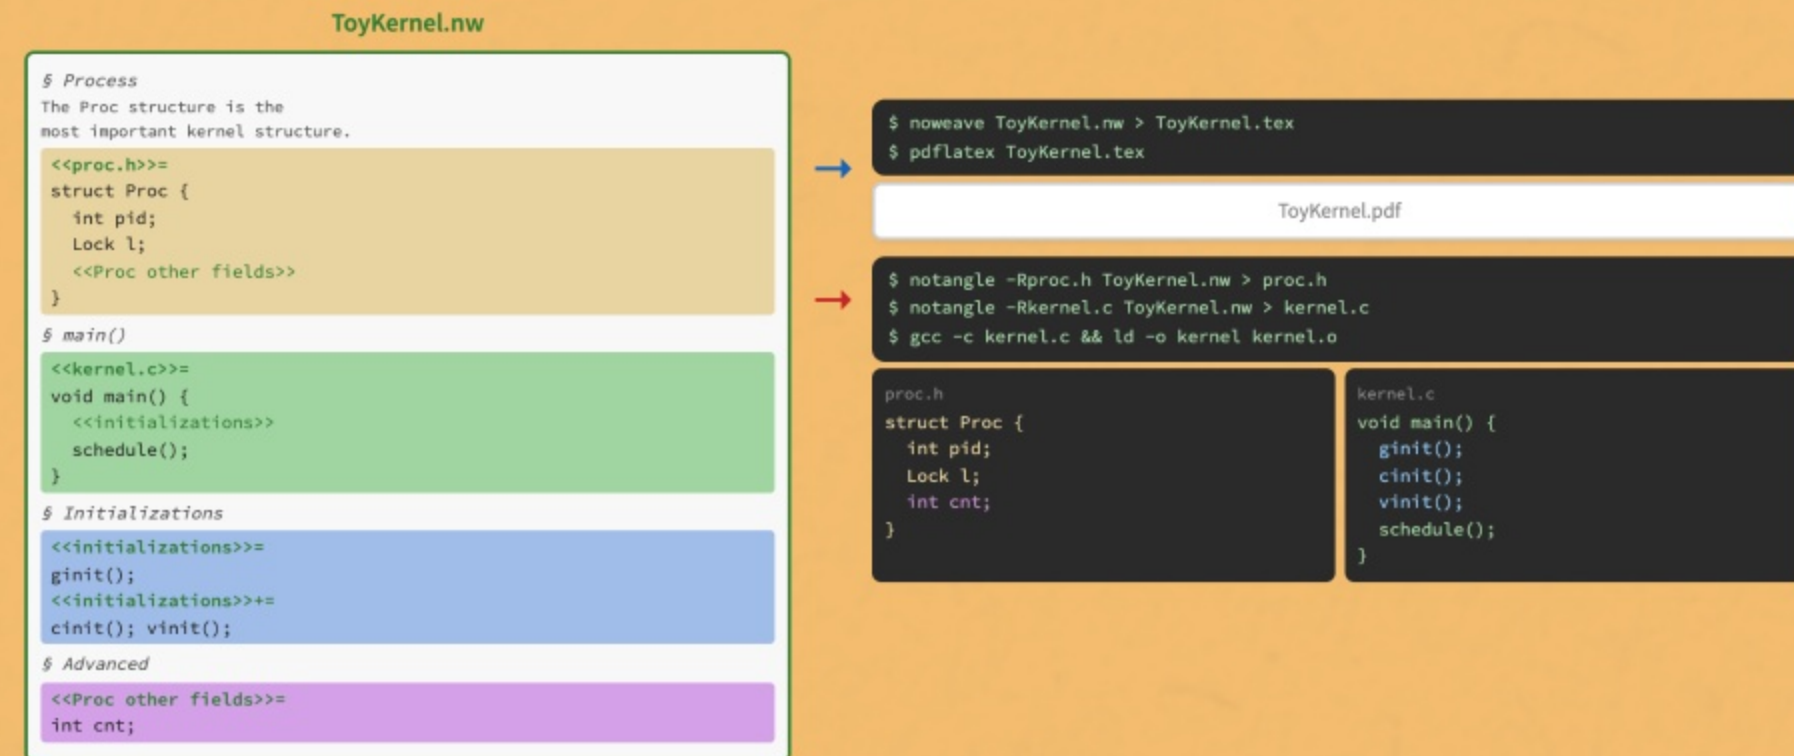
```
<<chunk-name>>= defines a chunk <<chunk-name>> references a chunk  
Chunks defined incrementally, across chapters
```

Literate programming 101

| ToyKernel.nw excerpt | L ^A T _E X output |
|--|--|
| <pre> \subsection{Process} The [[Proc]] structure is the most important kernel structure. <<proc.h>>= struct Proc { int pid; Lock l; <<[[Proc]] other fields>> } @ \subsection{[[main()]]} The kernel entry point is [[main()]] which after some initializations calls the most important function: [[schedule()]]. <<kernel.c>>= void main() { <<initializations>> schedule(); // never reached } @ \subsection{Initializations} This initializes all the globals: <<initializations>>= ginit(); @ Here are the remaining initializations: <<initializations>>= cinit(); vinit(); @ \subsection{Advanced features} Some less important [[Proc]] fields: <<[[Proc]] other fields>>= int cnt; @ </pre> | <pre> 3.1 Process The Proc structure is the most important ker- nel structure. (proc.h 6a)≡ struct Proc { int pid; Lock l; (Proc other fields 6e) } 3.2 main() The kernel entry point is main() which after some initializations calls the most important function: schedule(). (kernel.c 6b)≡ void main() { (initializations 6c) schedule(); // never reached } 3.3 Initializations This initializes all the globals: (initializations 6c)≡ ginit(); Here are the remaining initializations: (initializations 6d)+≡ cinit(); vinit(); 3.4 Advanced features Some less important Proc fields: (Proc other fields 6e)≡ int cnt; </pre> |

Table 2: Noweb Source Example and Rendered Output.

Literate Programming in Practice



| ToyKernel.nw excerpt | L^AT_EX output |
|--|--|
| <pre> \subsection{Process} The [[Proc]] structure is the most important kernel structure. <<proc.h>>= struct Proc { int pid; Lock l; <<[[Proc]] other fields>> } @ </pre> | <p>3.1 Proc</p> <p>The Proc stru nel structure.</p> <pre> <proc.h 6a>≡ struct Proc int pid; Lock l; <Proc other } </pre> |
| <pre> \subsection{[[main()]]} The kernel entry point is [[main()]] which after some initializations calls the most important function: [[schedule()]]. </pre> | <p>3.2 main()</p> <p>The kernel en some initializa function: sche</p> |

ToyKernel.nw excerpt

```
\subsection{Process}
```

The `[[Proc]]` structure is the most important kernel structure.

```
<<proc.h>>=  
struct Proc {  
    int pid;  
    Lock l;  
    <<[[Proc]] other fields>>  
}
```

@

```
\subsection{[[main()]]}
```

The kernel entry point is `[[main()]]` which after some initializations calls the most important function: `[[schedule()]]`.

```
<<kernel.c>>=  
void main() {  
    <<initializations>>  
    schedule();  
    // never reached  
}
```

L^AT_EX output

3.1 Process

The Proc structure is the most important kernel structure.

```
<proc.h 6a>≡  
    struct Proc {  
        int pid;  
        Lock l;  
        <Proc other fields 6e>  
    }
```

3.2 main()

The kernel entry point is `main()` which after some initializations calls the most important function: `schedule()`.

```
<kernel.c 6b>≡  
    void main() {  
        <initializations 6c>  
        schedule();  
        // never reached  
    }
```

```
<<kernel.c>>=
void main() {
    <<initializations>>
    schedule();
    // never reached
}
```

```
@
\subsection{Initializations}
This initializes all the globals:
```

```
<<initializations>>=
ginit();
@
Here are the remaining initializations:
```

```
<<initializations>>=
cinit();
vinit();
@
```

```
\subsection{Advanced features}
Some less important [[Proc]] fields:
```

```
<<[[Proc]] other fields>>=
int cnt;
@
```

```
<kernel.c 6b>≡
void main() {
    <initializations 6c>
    schedule();
    // never reached
}
```

3.3 Initializations

This initializes all the globals:

```
<initializations 6c>≡
ginit();
```

Here are the remaining initializations:

```
<initializations 6d>+≡
cinit();
vinit();
```

3.4 Advanced features

Some less important Proc fields:

```
<Proc other fields 6e>≡
int cnt;
```

Literate Programming in Practice

ToyKernel.nw

\$ Process

The Proc structure is the most important kernel structure.

```
<<proc.h>>=  
struct Proc {  
    int pid;  
    Lock l;  
    <<Proc other fields>>  
}
```

\$ main()

```
<<kernel.c>>=  
void main() {  
    <<initializations>>  
    schedule();  
}
```

\$ Initializations

```
<<initializations>>=  
ginit();  
<<initializations>>+=  
cinit(); vinit();
```

\$ Advanced

```
<<Proc other fields>>=  
int cnt;
```

`$ noweave ToyKernel.nw > ToyKernel.tex`
`$ pdflatex ToyKernel.tex`

ToyKernel.pdf

`$ notangle -Rproc.h ToyKernel.nw > proc.h`
`$ notangle -Rkernel.c ToyKernel.nw > kernel.c`
`$ gcc -c kernel.c && ld -o kernel kernel.o`

```
proc.h  
struct Proc {  
    int pid;  
    Lock l;  
    int cnt;  
}
```

```
kernel.c  
void main() {  
    ginit();  
    cinit();  
    vinit();  
    schedule();  
}
```

Why Books?

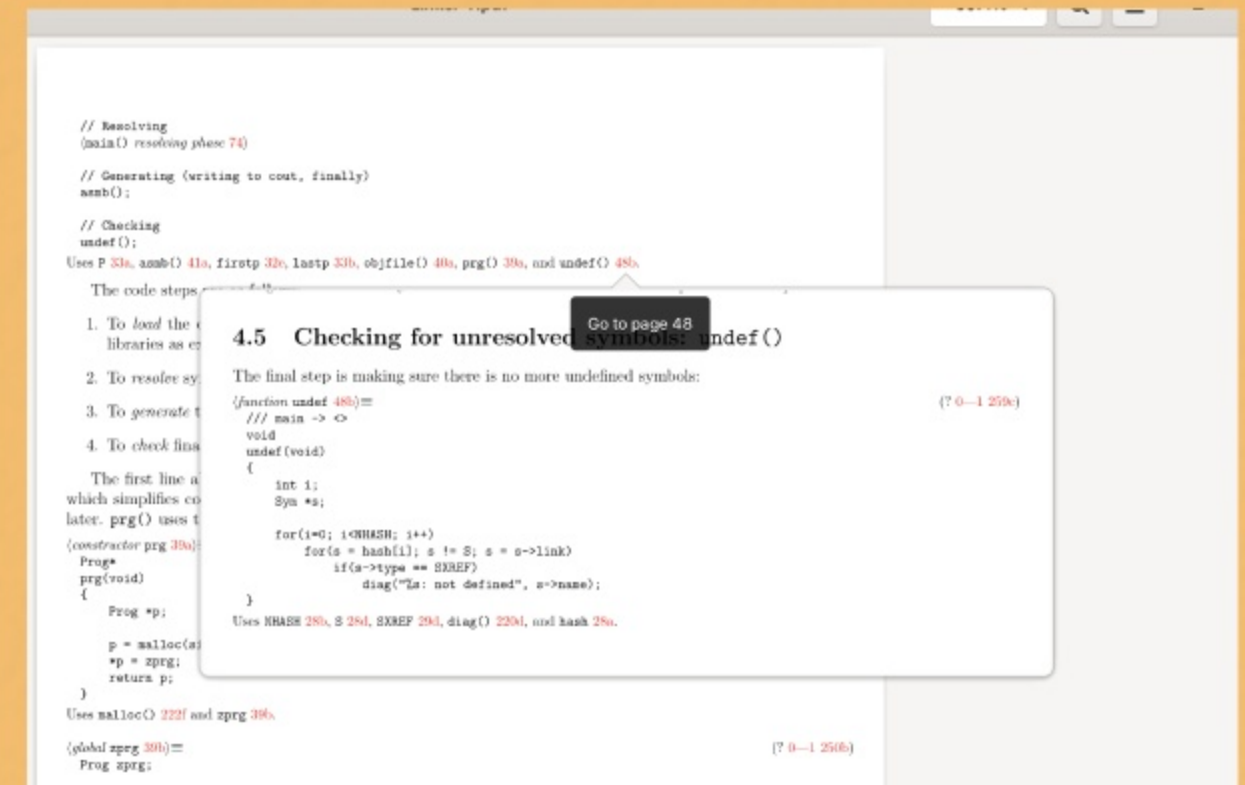
"Are books still the best medium in 2026?"

Books: teaching for **~2000 years**. Survived radio, TV, internet, MOOCs.

- **IDE:** compiler order, no narrative
- **Videos:** can't reference, can't search
- **AI explanation:** no coherent narrative

Syncweb brings IDE features INTO the book

Narrative of a book + navigation of an IDE



```
// Resolving
(main() resolving phase 74)

// Generating (writing to cout, finally)
asmb();

// Checking
undef();
```

Uses P 33a, asmb() 41a, firstp 32e, lastp 33b, objfile() 40a, prg() 39a, and undef() 48b.

The code steps

1. To load the libraries as ex
2. To resolve sy
3. To generate t
4. To check fina

The first line a which simplifies co later. prg() uses t

```
(constructor prg 39a):
Prog*
prg(void)
{
    Prog *p;

    p = malloc(si
    *p = zprg;
    return p;
}
```

Uses malloc() 222f and zprg 39b.

```
(global zprg 39b)≡
Prog zprg;
```

Go to page 48

4.5 Checking for unresolved symbols: undef()

The final step is making sure there is no more undefined symbols:

```
(function undef 48b)≡
/// main -> <>
void
undef(void)
{
    int i;
    Sym *s;

    for(i=0; i<NHASH; i++)
        for(s = hash[i]; s != S; s = s->link)
            if(s->type == SXREF)
                diag("%s: not defined", s->name);
}
```

Uses NHASH 28b, S 28d, SXREF 29d, diag() 220d, and hash 28a.

(? 0—1 259c)

(? 0—1 250b)

Transfer — Learn Here, Apply Everywhere

Understanding one small elegant OS gives deep intuition about **Linux**, **macOS**, and even Windows.

Plan 9 ideas already everywhere: **UTF-8**, **/proc**, Linux **namespaces** (=Docker), **Go**, **9P** in WSL2

Not just theory — you can **use** the tools:

plan9port: `acme`, `sam`, `rc`, `grep`, `sed`, `awk` on Linux/macOS

goken9cc: Plan 9 compilers that build Linux/macOS/Windows programs

Use them in your own environment, outside Plan 9.

The same tools **AI coding agents run hundreds of times a day**.

Same concepts, 100x more clearly. And you can use them for real.

Principia Softwarica

Plan 9 Code Explained

Yoann Padioleau

IWP9



The Best Engineers Understand the Stack

Facebook: linker slow, but benchmark for everyone
→ most engineers: stack, accepted it, black box

We had to wait for [Ian Lance Taylor](#) (Google, 2006) to write `glibc` → a new linker for GNU Binutils.

Not a young hotshot. Old school Unix/GCC hacker (since 1996).
Also part of the Go team at Google.
The kind of engineer who understands the full stack.
Understand deeply = build better

The Education Gap

"What happens when you type `ls` in a terminal window?"

[Keith Adams](#) → colleague at Facebook, later Chief Architect at Slack → used this as his interview question.

Simple question. The answer involves: the shell, the C library, the kernel, the graphics stack, the windowing system.
Most engineers can't answer.

"Full stack" today = React + Node + cloud.
REAL full stack = compiler, linker, kernel, graphics

The AI Era Makes This More Important

We write less code. We need **MORE** understanding.

AI generates code → but who evaluates it?
Who debugs the segfault? The linker error?

AI tools like `CLAUDE` Code USE the classic commands:
`grep`, `ls`, `diff`, `awk`, `gcc`, `ls` ...
The AI is a power user of the very programs Principia explains.

Writing code = 20%. Understanding = 80%.
AI handles the 20%. The 80% becomes everything.

But the Stack is Enormous

| | |
|--------------|-----------|
| Linux kernel | ~40M LOC |
| glibc | ~15M LOC |
| bash | ~3.5M LOC |
| log server | ~100M LOC |
| stern alone | ~80M LOC |

Each component: more code than ALL of Plan 9

Plan 9 - 100k Lines Code

| | |
|-------------|----------|
| Kernel | 100k LOC |
| glibc | 15M LOC |
| bash | 3.5M LOC |
| log server | 100M LOC |
| stern alone | 80M LOC |

Principia Software

A Series of Books

| Category | Book | Program(s) | LOC | LOE | Pages |
|-----------------------|------------------|---|---------------|--------------|-------------|
| Core system | Kernel | 9pi | 35000 | 3200 | 732 |
| | Core libraries | libc libregexp libthread | 19000 | 1600 | 438 |
| | Shell | rc | 6500 | 1700 | 166 |
| Development toolchain | C compiler | 5c libcc | 18500 | 1900 | 471 |
| | Assembler | 5a | 3600 | 4400 | 176 |
| | Linker | 5l | 7500 | 5400 | 296 |
| Developer tools | Editor | ed | 1600 | 200 | 45 |
| | Build system | mk | 4350 | 4050 | 197 |
| | Debuggers | db acid strace libmach | 13100 | 1000 | 321 |
| | Profilers | prof time kprof stats iostats | 3900 | 350 | 102 |
| Graphics | Graphics stack | /dev/draw libdraw libmemdraw libmemlayer | 18500 | 3400 | 507 |
| | Windowing system | rio libframe libcomplete libplumb | 8800 | 4000 | 289 |
| Networking | Network stack | /dev/net libip lib9p | 18300 | 4800 | 457 |
| Misc | CLI utilities | cat ls grep sed diff tar gzip ... | 23900 | 650 | 493 |
| Total: | | | 182550 | 36650 | 4690 |

Table 1: Principia Softwarica Books and Their Statistics.
(LOC = lines of code; LOE = lines of explanation; Pages = number of typeset pages)

A New Plan 9 Fork

```

Why a New Fork?
https://github.com/aryx/principia-softwarica

Repository structure
Screenshot: Repository structure diagram showing folders like assemblers/, kernel/, mkfiles/, etc.

Screenshot: Repository structure diagram showing folders like assemblers/, kernel/, mkfiles/, etc.

Screenshot: Repository structure diagram showing folders like assemblers/, kernel/, mkfiles/, etc.

-- assemblers/
|   |-- 5a/
|   |-- 5a/
|   |-- data2s.c
|   |-- Assembler.nw
|   |-- mkfile
-- builders/
|   |-- Builder.nw
|   |-- mk/
|   |-- mkfile
-- compilers/
|   |-- 5c/
|   |-- 5c/
|   |-- 5c/
|   |-- cc/
|   |-- Compiler.nw
|   |-- cpp/
|   |-- mkfile
-- debuggers/
|   |-- acid/
|   |-- db/
|   |-- Debugger.nw
|   |-- libmach/
|   |-- mkfile
-- Dockerfile
-- docs/
|   |-- articles/
|   |-- iw9/
-- dosdisk.img
-- editors/
|   |-- ed/
|   |-- mkfile
-- include/
|   |-- libc.h
|   |-- ...
|   |-- u.h
-- linkers/
|   |-- 5l/
|   |-- 5l/
|   |-- Linker.nw
|   |-- mkfile
-- readme.txt
-- windows/
|   |-- libcomplete/
|   |-- libplumb/
|   |-- mkfile
|   |-- rio/
|   |-- Windows.nw
|   |-- arch/
|   |-- concurrency/
|   |-- console/
|   |-- devices/
|   |-- files/
|   |-- filesystems/
|   |-- init/
|   |-- memory/
|   |-- Kernel.nw
|   |-- mkfile
|   |-- network/
|   |-- processes/
|   |-- syscalls/
|   |-- lib_core/
|   |-- libbio/
|   |-- libc/
|   |-- Libcore.nw
|   |-- libthread/
|   |-- mkfile
|   |-- lib_graphics/
|   |-- Graphics.nw
|   |-- libdraw/
|   |-- libing/
|   |-- libmemdraw/
|   |-- libmemlayer/
|   |-- mkfile
|   |-- lib_networking/
|   |-- lib9p/
|   |-- libip/
|   |-- mkfile
|   |-- lib_strings/
|   |-- libflate/
|   |-- libregexp/
|   |-- libstring/
|   |-- mkfile
|   |-- Makefile
|   |-- mkconfig.pc
|   |-- mkconfig.pi
|   |-- mkfile
|   |-- mkfile-host-Cygwin
|   |-- mkfile-host-Linux
|   |-- mkfile-host-macOS
|   |-- mkfile-host-Plan9
|   |-- MISC/
|   |-- pc/
|   |-- pi/
|   |-- 386/
|   |-- arm/
|   |-- arm/
|   |-- lib/
|   |-- rc/
|   |-- tests/
|   |-- usr/

```

Figure 1: <https://github.com/aryx/principia-softwarica> Layout.
(as generated by `tree -L 2 -F` and spread in 3 columns)

A New Way to Build

Goken Dockerfile excerpt

```

# Build goken on Ubuntu using
# gcc/binutils (and mk/rc)
FROM ubuntu:24.04

# Setup a basic C dev environment
RUN apt-get install -y gcc libc6-dev \
    byacc

# Now let's build from source
WORKDIR /src
COPY . .

# Small shell script (not GNU autoconf)
# to detect arch and generate
# ./mkconfig
RUN ./configure

# The script below obviously builds
# 'mk' (without needing mk) but also:
# - 'rc', which is called by 'mk'
# - 'ed', for the mkenam script
# run during the build
RUN ./scripts/build-mk.sh

# copy ./ROOT/<arch>/bin/{mk,rc,ed}
# to ./bin/
RUN ./scripts/promote-mk.sh

# make mk and rc accessible
ENV PATH="/src/bin:${PATH}"
ENV MKSHELL="/src/bin/rc"

# Let's build goken (using mk/rc built
# in the previous step)
RUN mk
RUN mk install

ENV PATH="/src/ROOT/amd64/bin:\
    /src/ROOT/arm64/bin:\
    ${PATH}"

```

Principia Softwarica Dockerfile

```

# Build principia on Ubuntu
# Linux (amd64 or arm64)
# for 386 (pc) and arm (pi)
FROM padator/goken

WORKDIR /principia
COPY . .

# 386
RUN cp mkconfig.pc mkconfig
RUN mk && mk install
RUN mk kernel

# arm
RUN cp mkconfig.pi mkconfig
RUN mk && mk install
RUN mk kernel

# VFAT tools
RUN apt-get install -y dosfstools \
    mtools

# making dosdisk.img
RUN dd if=/dev/zero of=tmp.img \
    bs=1M count=512
RUN mks.vfat tmp.img
RUN mcopy -i tmp.img -s -o \
    ROOT/* : :
RUN cat MISC/pc/bootsector tmp.img \
    > dosdisk.img
RUN rm -f tmp.img

```

Table 3: Principia Softwarica Build and Packaging Process.

A New Plan 9 Fork

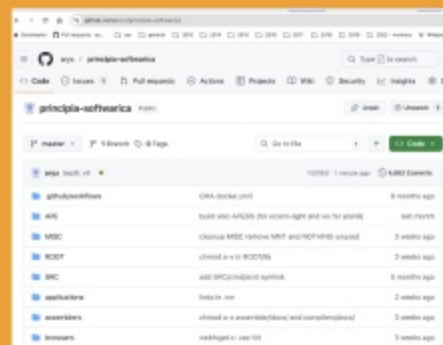
Why a New Fork?

Plan 9 = research OS. Principia = teaching OS. Different needs.



Across all books: introduced `globals.c`, `reemulate` .h files, renamed for clarity, removed dead code, fixed bugs.
Result: code extensively reorganized → fork needed

<https://github.com/aryx/principia-softwareica>



```
.
|-- assemblers/
|   |-- 5a/
|   |-- 8a/
|   |-- data2s.c
|   |-- Assembler.nw
|   '-- mkfile
-- builders/
|   |-- Builder.nw
|   |-- mk/
|   '-- mkfile
-- compilers/
|   |-- 5c/
|   |-- 8c/
|   |-- cc/
|   |-- Compiler.nw
|   |-- cpp/
|   '-- mkfile
-- debuggers/
|   |-- acid/
|   |-- db/
|   |-- Debugger.nw
|   |-- libmach/
|   '-- mkfile
-- Dockerfile
|-- kernel/
|   |-- arch/
|   |-- concurrency/
|   |-- console/
|   |-- devices/
|   |-- files/
|   |-- filesystems/
|   |-- init/
|   |-- memory/
|   |-- Kernel.nw
|   |-- mkfile
|   |-- network/
|   |-- processes/
|   '-- syscalls/
-- lib_core/
|   |-- libbio/
|   |-- libc/
|   |-- Libcore.nw
|   |-- libthread/
|   '-- mkfile
-- lib_graphics/
|   |-- Graphics.nw
|   |-- libdraw/
|   |-- libimg/
|-- mkfiles/
|   |-- 386/
|   |-- arm/
|   |-- mkdirs
|   |-- mkfile.proto
|   |-- mklib
|   '-- mkone
-- mkfile-target-pc
-- mkfile-target-pi
-- networking/
|   |-- arp/
|   |-- dhcp/
|   |-- ftp/
|   |-- http/
|   |-- ip/
|   |-- mkfile
|   |-- ndb/
|   |-- Network.nw
|   |-- snoopy/
|   '-- telnet/
-- profilers/
|   |-- iostats/
|   |-- mkfile
|   '-- Profiler.nw
```

Why a New Fork?

Plan 9 = research OS. Principia = **teaching OS**. Different needs.

Plan 9 4th edition

sys/src/9/port/

proc.c – everything in one file:

schedinit() sleep() procinit() sched()
wakeup() pexit() qlock

page.c

devcons.c

~~dead code...~~

sys/src/9/pc/

trap.c

sys/src/9/bcm/

trap.c (same name, diff arch)

Flat files, mixed concerns, dead code

→
split
rename
clean
reorganize

Principia

Kernel.nw (one book, chapters by concern)

Ch. Processes: procinit() pexit()

Ch. Concurrency: sleep() wakeup()

Ch. Scheduling: schedinit() sched()

Ch. Memory: page functions

Ch. Interrupts: trap() port + arch unified

Ch. Advanced: complex details deferred

kernel/ processes/ proc.c sysproc.c · concurrency/ qlock.c taslock.c · memory/
page.c segment.c · interrupts/ · scheduling/ ...

Files reorganized into topic dirs too

Across all books: introduced `globals.c`, intermediate `.h` files, renamed for clarity, removed dead code, fixed bugs

Result: code extensively reorganized → **fork needed**

<https://github.com/aryx/principia-softwarica>

The screenshot shows the GitHub repository page for 'principia-softwarica' by user 'aryx'. The repository is public and has 1 branch (master) and 0 tags. The commit history shows 4,662 commits, with the latest commit being 113785f, made 1 minute ago. The repository contains several folders, each with a commit message and a timestamp:

| Folder | Commit Message | Timestamp |
|-------------------|--|--------------|
| .github/workflows | GHA docker.yml! | 9 months ago |
| APE | build also APE/lib (for ocaml-light and xix for plan9) | last month |
| MISC | cleanup MISC remove MNT and NOTHING unused | 3 weeks ago |
| ROOT | chmod a-x in ROOT/lib | 3 weeks ago |
| SRC | add SRC/cmd/acid symlink | 5 months ago |
| applications | links in .nw | 2 weeks ago |
| assemblers | chmod a-x assembler/docs/ and compilers/docs/ | 3 weeks ago |
| browsers | webfsget.c: use fdt | 3 weeks ago |

```

.
|-- assemblers/      |-- kernel/          |-- mkfiles/
|  |-- 5a/           |  |-- arch/         |  |-- 386/
|  |-- 8a/           |  |-- concurrency/ |  |-- arm/
|  |-- data2s.c      |  |-- console/     |  |-- mkdirs
|  |-- Assembler.nw |  |-- devices/     |  |-- mkfile.proto
|  '-- mkfile        |  |-- files/       |  |-- mklib
|-- builders/       |  |-- filesystems/ |  '-- mkone
|  |-- Builder.nw   |  |-- init/        |-- mkfile-target-pc
|  |-- mk/          |  |-- memory/     |-- mkfile-target-pi
|  '-- mkfile       |  |-- Kernel.nw  |-- networking/
|-- compilers/     |  |-- mkfile      |  |-- arp/
|  |-- 5c/          |  |-- network/    |  |-- dhcp/
|  |-- 8c/          |  |-- processes/  |  |-- ftp/
|  |-- cc/          |  '-- syscalls/   |  |-- http/
|  |-- Compiler.nw |  |-- lib_core/   |  |-- ip/
|  |-- cpp/         |  |-- libbio/     |  |-- mkfile
|  '-- mkfile       |  |-- libc/       |  |-- ndb/
|-- debuggers/     |  |-- Libcore.nw |  |-- Network.nw
|  |-- acid/        |  |-- libthread/  |  |-- snoopy/
|  |-- db/          |  '-- mkfile      |  '-- telnet/
|  |-- Debugger.nw |  |-- lib_graphics/ |-- profilers/

```

Repository Structure

- `assemblers/5a/` not `sys/src/cmd/5a`
- Each dir: `.nw` literate program + generated code
- `kernel/`: `concurrency/` `memory/` `processes/` ...
- `mkconfig.pc` (x86/QEMU) + `mkconfig.pi` (ARM/RPi)

Good strategy to find bugs: try to explain the code in a book.

Syncweb: Bidirectional LP

Classic Noweb: `.nw` \longrightarrow `code` (one way) Syncweb: `.nw` \longleftrightarrow `code` (both ways)

Marks + md5 checksums \rightarrow detect which side changed

Automatic merge or conflict detection

Solves biggest LP complaint

github.com/aryx/syncweb

Principia Softwarica

Plan 9 Code Explained

Yoann Padioleau

IWP9



The Best Engineers Understand the Stack

Facebook: linker slow, but benchmark for everyone
→ most engineers: stack, accepted it, black box

We had to wait for [Ian Lance Taylor](#) (Google, 2006) to write `glibc` → a new linker for GNU Binutils.

Not a young hotshot. Old school Unix/GCC hacker (since 1996).
Also part of the Go team at Google.
The kind of engineer who understands the full stack.
Understand deeply = build better

The Education Gap

"What happens when you type `ls` in a terminal window?"

[Keith Adcox](#) → colleague at Facebook, later Chief Architect at Slack → used this as his interview question.

Simple question. The answer involves: the shell, the C library, the kernel, the graphics stack, the windowing system.
Most engineers can't answer.

"Full stack" today = React + Node + cloud.
REAL full stack = compiler, linker, kernel, graphics

The AI Era Makes This More Important

We write less code. We need **MORE** understanding.

AI generates code → but who evaluates it?
Who debugs the segfault? The linker error?

AI tools like `CLAUDE` Code USE the classic commands:
`grep`, `ls`, `diff`, `awk`, `gcc`, `ls` ...
The AI is a power user of the very programs Principia explains.

Writing code = 20%. Understanding = 80%.
AI handles the 20%. The 80% becomes everything.

But the Stack is Enormous

| | |
|--------------|-----------|
| Linux kernel | ~40M LOC |
| glibc | ~15M LOC |
| bash | ~3.5M LOC |
| log server | ~100M LOC |
| stern alone | ~80M LOC |

Each component: more code than ALL of Plan 9

Plan 9 - 100k Lines Code

| | |
|-------------|-----------|
| Kernel | ~100k LOC |
| glibc | ~100k LOC |
| bash | ~100k LOC |
| log server | ~100k LOC |
| stern alone | ~100k LOC |

Principia Software

A Series of Books

| Category | Book | Program(s) | LOC | LOE | Pages |
|-----------------------|------------------|---|---------------|--------------|-------------|
| Core system | Kernel | 9pi | 35000 | 3200 | 732 |
| | Core libraries | libc libregexp libthread | 19000 | 1600 | 438 |
| | Shell | rc | 6500 | 1700 | 166 |
| Development toolchain | C compiler | 5c libcc | 18500 | 1900 | 471 |
| | Assembler | 5a | 3600 | 4400 | 176 |
| | Linker | 5l | 7500 | 5400 | 296 |
| Developer tools | Editor | ed | 1600 | 200 | 45 |
| | Build system | mk | 4350 | 4050 | 197 |
| | Debuggers | db acid strace libmach | 13100 | 1000 | 321 |
| Graphics | Graphics stack | prof time kprof stats iostats | 3900 | 350 | 102 |
| | | /dev/draw | 18500 | 3400 | 507 |
| | | libdraw libmemdraw libmemlayer | | | |
| Networking | Windowing system | rio libframe libcomplete libplumb | 8800 | 4000 | 289 |
| | | /dev/net | 18300 | 4800 | 457 |
| | | libip lib9p | | | |
| Misc | CLI utilities | cat ls grep sed diff tar gzip ... | 23900 | 650 | 493 |
| Total: | | | 182550 | 36650 | 4690 |

Table 1: Principia Softwarica Books and Their Statistics.
(LOC = lines of code; LOE = lines of explanation; Pages = number of typeset pages)

A New Plan 9 Fork

```

-- assemblers/      -- kernel/          -- mkfiles/
-- 5a/              -- arch/            -- 386/
-- 5a/              -- concurrency/    -- arm/
-- data2s.c         -- console/         -- makedirs
-- Assembler.nw    -- devices/         -- mkfile.proto
-- mkfile           -- files/           -- mklib
-- builders/       -- filesystems/    -- mkone
-- Builder.nw       -- init/            -- mkfile-target-pc
-- mk/              -- memory/         -- mkfile-target-pi
-- mkfile           -- Kernel.nw       -- networking/
-- compilers/      -- mkfile           -- arp/
-- 5c/              -- network/         -- dhcp/
-- 5c/              -- processes/      -- ftp/
-- cc/              -- syscalls/       -- http/
-- Compiler.nw     -- lib_core/       -- ip/
-- cpp/            -- libbio/          -- mkfile
-- mkfile           -- libc/            -- ndb/
-- debuggers/     -- libcore.nw      -- Network.nw
-- acid/           -- libthread/      -- snoopy/
-- db/             -- mkfile          -- telnet/
-- Debugger.nw    -- lib_graphics/   -- profilers/
-- libmach/        -- Graphics.nw    -- iostats/
-- mkfile          -- libdraw/        -- mkfile
-- Dockerfile      -- libing/         -- Profiler.nw
-- docs/           -- libmemdraw/    -- shells/
-- articles/       -- libmenlayer/   -- mkfile
-- iw9p/           -- mkfile          -- rc/
-- dosdisk.img     -- lib_networking/ -- Shell.nw
-- editors/        -- lib9p/          -- utilities/
-- ed/             -- libip/          -- archive/
-- mkfile          -- mkfile          -- byte/
-- include/        -- lib_strings/    -- calc/
-- lib.h           -- libflate/       -- compare/
-- ...            -- libregexp/     -- files/
-- u.h             -- libstring/     -- mkfile
-- linkers/        -- mkfile          -- pipe/
-- 5l/             -- Makefile        -- process/
-- 5l/             -- mkconfig.pc    -- text/
-- Linker.nw       -- mkconfig.pi    -- time/
-- mkfile          -- mkfile          -- Utilities.nw
-- readme.txt     -- mkfile-host-Cygwin
-- windows/        -- mkfile-host-Linux
-- libcomplete/    -- mkfile-host-macOS
-- libplumb/       -- mkfile-host-Plan9
-- mkfile          -- MISC/
-- rio/            -- pc/
-- Windows.nw     -- pi/
  
```

Figure 1: <https://github.com/aryx/principia-softwarica> Layout.
(as generated by `tree -L 2 -F` and spread in 3 columns)

A New Way to Build

Goken Dockerfile excerpt

```

# Build goken on Ubuntu using
# gcc/binutils (and mk/rc)
FROM ubuntu:24.04

# Setup a basic C dev environment
RUN apt-get install -y gcc libc6-dev \
    byacc

# Now let's build from source
WORKDIR /src
COPY . .

# Small shell script (not GNU autoconf)
# to detect arch and generate
# ./mkconfig
RUN ./configure

# The script below obviously builds
# 'mk' (without needing mk) but also:
# - 'rc', which is called by 'mk'
# - 'ed', for the mkenam script
# run during the build
RUN ./scripts/build-mk.sh

# copy ./ROOT/<arch>/bin/{mk,rc,ed}
# to ./bin/
RUN ./scripts/promote-mk.sh

# make mk and rc accessible
ENV PATH="/src/bin:${PATH}"
ENV MKSHELL="/src/bin/rc"

# Let's build goken (using mk/rc built
# in the previous step)
RUN mk
RUN mk install

ENV PATH="/src/ROOT/amd64/bin:\
    /src/ROOT/arm64/bin:\
    ${PATH}"
  
```

Principia Softwarica Dockerfile

```

# Build principia on Ubuntu
# Linux (amd64 or arm64)
# for 386 (pc) and arm (pi)
FROM padator/goken

WORKDIR /principia
COPY . .

# 386
RUN cp mkconfig.pc mkconfig
RUN mk && mk install
RUN mk kernel

# arm
RUN cp mkconfig.pi mkconfig
RUN mk && mk install
RUN mk kernel

# VFAT tools
RUN apt-get install -y dosfstools \
    mtools

# making dosdisk.img
RUN dd if=/dev/zero of=tmp.img \
    bs=1M count=512
RUN mks.vfat tmp.img
RUN mcopy -i tmp.img -s -o \
    ROOT/* : :
RUN cat MISC/pc/bootsector tmp.img \
    > dosdisk.img
RUN rm -f tmp.img
  
```

Table 3: Principia Softwarica Build and Packaging Process.

A New Way to Build

The Build Problem

Teaching context: students want to modify code, rebuild, and boot
— from Linux, macOS, Windows, in their own editor or IDE.

Problem 1: Plan 9 uses a special C dialect with extensions

Anonymous structs, different calling conventions, custom assembler
→ gcc / clang can't compile it as-is

Problem 2: Existing options (Bfront, Blegacy) have

non-trivial build processes, packaging not straightforward

Prior art: **Inferno** (Vita Nuova), **plan9port** (Russ Cox), **kence** (Blegacy)

Each made steps toward cross-compiling Plan 9 from Linux/OS

goken9cc: Plan 9 compilers, built with gcc/clang

Cross-compile from Linux/macOS/Windows → boot with QEMU or Raspberry Pi

github.com/aryx/goken9cc

5 Commands: Clone, Build, Run

```
$ git clone https://github.com/aryx/goken9cc $ cd goken9cc && docker build -t "padator/goken" . $ cd ../principia-software $ docker build -t "principia" . $ qemu-system-i386 -m 4 -m 512 \ -kernel $qemu -hda dosdisk.img
```

Irony: Docker relies on Linux namespaces — inspired by Plan 9

github.com/aryx/goken9cc

Live Demo: Modify, Rebuild, Boot

```
BEFORE: $ strace ls - open = syscall 4
CHANGE: $ vi lib_core/libc/syscall/sys.h #define OPEN 0 - #define OPEN 42
RESULT: $ time mk && mk install && mk kernel < 10 seconds
BOOT: $ qemu-system-i386 -kernel $qemu -hda dosdisk.img $ strace ls - open = syscall 42
```

"Rebuilt the entire OS in under a minute. Try that with Linux."

That [ls](#) went through shell, libc, kernel, graphics, windowing. Every layer — explained in a Principia book.

Goken Dockerfile excerpt

```
# Build goken on Ubuntu using
# gcc/binutils (and mk/rc)
FROM ubuntu:24.04

# Setup a basic C dev environment
RUN apt-get install -y gcc libc6-dev \
    byacc

# Now let's build from source
WORKDIR /src
COPY . .

# Small shell script (not GNU autoconf)
# to detect arch and generate
# ./mkconfig
RUN ./configure
```

```
# The script below obviously builds
# 'mk' (without needing mk) but also:
```

Principia Softwarica Dockerfile

```
# Build principia on Ubuntu
# Linux (amd64 or arm64)
# for 386 (pc) and arm (pi)
FROM padator/goken

WORKDIR /principia
COPY . .

# 386
RUN cp mkconfig.pc mkconfig
RUN mk && mk install
RUN mk kernel

# arm
RUN cp mkconfig.pi mkconfig
RUN mk && mk install
RUN mk kernel
```

```
# VFAT tools
RUN apt-get install -v dosfstools \
```

The Build Problem

Teaching context: students want to **modify** code, **rebuild**, and **boot** — from Linux, macOS, Windows, in their own editor or IDE.

Problem 1: Plan 9 uses a special C dialect with extensions

Anonymous structs, different calling conventions, custom assembler

→ `gcc` / `clang` can't compile it as-is

Problem 2: Existing options (9front, 9legacy) have

non-trivial build procedures, packaging not straightforward

Prior art: **Inferno** (Vita Nuova), **plan9port** (Russ Cox), **kenc** (9legacy)

Each made steps toward cross-compiling Plan 9 code from Unix/Linux

goken9cc: Plan 9 compilers, built with gcc/clang

Cross-compile from Linux/macOS/Windows → boot with QEMU or Raspberry Pi

github.com/aryx/goken9cc

Goken Dockerfile excerpt

```
# Build goken on Ubuntu using
# gcc/binutils (and mk/rc)
FROM ubuntu:24.04

# Setup a basic C dev environment
RUN apt-get install -y gcc libc6-dev \
    byacc

# Now let's build from source
WORKDIR /src
COPY . .
```

Want to **modify** code, **rebuild**, and **boot**
flows, in their own editor or IDE.

Special C dialect with extensions
including conventions, custom assembler
and it as-is

(9front, 9legacy) have
packaging not straightforward

(a), **plan9port** (Russ Cox), **kenc** (9legacy)
compiling Plan 9 code from Unix/Linux

compilers, built with gcc/clang

OS/Windows → boot with QEMU or Raspberry Pi

, Run

```
/aryx/goken9cc $ cd goken9cc && docker build -t  
cipia-softwarica $ docker build -t "principia" . $  
\ -kernel 9qemu -hda dosdisk.img
```

aces — inspired by Plan 9

ebuild, Boot

```
yscall 6
```

```
scall/sys.h #define OPEN 6 - #define OPEN 42
```

```
ll && mk kernel < 10 seconds
```

```
el 9qemu -hda dosdisk.img $ strace ls -> open =
```

nder a minute. Try that with Linux."

nel, graphics, windowing. Every layer — explained in a

```
# Small shell script (not GNU autoconf)
# to detect arch and generate
# ./mkconfig
RUN ./configure
```

```
# The script below obviously builds
# 'mk' (without needing mk) but also:
# - 'rc', which is called by 'mk'
# - 'ed', for the mkenam script
# run during the build
RUN ./scripts/build-mk.sh
```

```
# copy ./ROOT/<arch>/bin/{mk,rc,ed}
# to ./bin/
RUN ./scripts/promote-mk.sh
```

```
RUN mk kerne
# arm
RUN cp mkcon
RUN mk && mk
RUN mk kerne
```

```
# VFAT tools
RUN apt-get
      mtools
```

```
# making dos
RUN dd if=/d
      bs=1
```

```
RUN mkfs.vfa
RUN mcopy -i
      ROOT/*
RUN cat MISC
```

```
# make mk and rc accessible
ENV PATH="/src/bin:${PATH}"
ENV MKSHELL="/src/bin/rc"
```

```
# Let's build goken (using mk/rc built
# in the previous step)
```

```
RUN mk
```

```
RUN mk install
```

```
ENV PATH="/src/ROOT/amd64/bin:\
        /src/ROOT/arm64/bin:\
        ${PATH}"
```

```
> dosdi
```

```
RUN rm -f tm
```

Table 3: Principia Softwarica Build and Packaging Pro

Principia Softwarica Dockerfile

```
# Build principia on Ubuntu  
# Linux (amd64 or arm64)  
# for 386 (pc) and arm (pi)  
FROM padator/goken
```

```
WORKDIR /principia  
COPY . .
```

```
# 386
```

```
RUN cp mkconfig.pc mkconfig
```

```
RUN mk && mk install
```

```
RUN mk kernel
```

```
# arm
```

```
RUN cp mkconfig.pi mkconfig
```

```
RUN mk && mk install
```

```
RUN mk kernel
```

VFAT tools

RUN apt-get install -y dosfstools \
mtools

making dosdisk.img

RUN dd if=/dev/zero of=tmp.img \
bs=1M count=512

RUN mkfs.vfat tmp.img

RUN mcopy -i tmp.img -s -o \
ROOT/* ::

RUN cat MISC/pc/bootsector tmp.img \
> dosdisk.img

RUN rm -f tmp.img

5 Commands: Clone, Build, Run

```
$ git clone https://github.com/aryx/goken9cc $ cd goken9cc && docker build -t "padator/goken" . $ cd ../principia-softwarica $ docker build -t "principia" . $ qemu-system-i386 -smp 4 -m 512 \ -kernel 9qemu -hda dosdisk.img
```

Irony: Docker relies on Linux namespaces — inspired by Plan 9

github.com/aryx/goken9cc

Live Demo: Modify, Rebuild, Boot

```
BEFORE: $ strace ls → open = syscall 6
```

```
CHANGE: $ vi lib_core/libc/9syscall/sys.h #define OPEN 6 → #define OPEN 42
```

```
REBUILD: $ time mk && mk install && mk kernel < 10 seconds
```

```
BOOT: $ qemu-system-i386 -kernel 9qemu -hda dosdisk.img $ strace ls → open =  
syscall 42
```

"Rebuilt the entire OS in under a minute. Try that with Linux."

That `ls` went through shell, libc, kernel, graphics, windowing. Every layer — explained in a Principia book.

Principia Softwarica

Plan 9 Code Explained

Yoann Padioleau

IWP9



The Best Engineers Understand the Stack

Facebook: linker slow, but benchmark for everyone
→ most engineers: stack, accepted it, black box

We had to wait for [Ian Lance Taylor](#) (Google, 2006) to write `glibc` → a new linker for GNU Binutils.

Not a young hotshot. Old school Unix/GCC hacker (since 1996).
Also part of the Go team at Google.
The kind of engineer who understands the full stack.
Understand deeply = build better

The Education Gap

"What happens when you type `ls` in a terminal window?"

[Keith Adcox](#) → colleague at Facebook, later Chief Architect at Slack → used this as his interview question.

Simple question. The answer involves: the shell, the C library, the kernel, the graphics stack, the windowing system.
Most engineers can't answer.

"Full stack" today = React + Node + cloud.
REAL full stack = compiler, linker, kernel, graphics

The AI Era Makes This More Important

We write less code. We need **MORE** understanding.

AI generates code → but who evaluates it?
Who debugs the segfault? The linker error?

AI tools like `CLAUDE` Code USE the classic commands:
`grep`, `ls`, `diff`, `awk`, `gcc`, `ls` ...
The AI is a power user of the very programs Principia explains.

Writing code = 20%. Understanding = 80%.
AI handles the 20%. The 80% becomes everything.

But the Stack is Enormous

| | |
|--------------|-----------|
| Linux kernel | ~40M LOC |
| glibc | ~15M LOC |
| bash | ~3.5M LOC |
| log server | ~100M LOC |
| stern alone | ~80M LOC |

Each component: more code than ALL of Plan 9

A Series of Books

| Category | Book | Program(s) | LOC | LOE | Pages |
|-----------------------|------------------|---|---------------|--------------|-------------|
| Core system | Kernel | 9pi | 35000 | 3200 | 732 |
| | Core libraries | libc libregexp libthread | 19000 | 1600 | 438 |
| | Shell | rc | 6500 | 1700 | 166 |
| Development toolchain | C compiler | 5c libcc | 18500 | 1900 | 471 |
| | Assembler | 5a | 3600 | 4400 | 176 |
| | Linker | 5l | 7500 | 5400 | 296 |
| Developer tools | Editor | ed | 1600 | 200 | 45 |
| | Build system | mk | 4350 | 4050 | 197 |
| | Debuggers | db acid strace libmach | 13100 | 1000 | 321 |
| Graphics | Graphics stack | prof time kprof stats iostats | 3900 | 350 | 102 |
| | | /dev/draw | 18500 | 3400 | 507 |
| | | libdraw libmemdraw libmemlayer | | | |
| Networking | Windowing system | rio libframe libcomplete libplumb | 8800 | 4000 | 289 |
| | | /dev/net | 18300 | 4800 | 457 |
| | | libip lib9p | | | |
| Misc | CLI utilities | cat ls grep sed diff tar gzip ... | 23900 | 650 | 493 |
| Total: | | | 182550 | 36650 | 4690 |

Table 1: Principia Softwarica Books and Their Statistics.
(LOC = lines of code; LOE = lines of explanation; Pages = number of typeset pages)

A New Plan 9 Fork

```

-- assemblers/
-- 5a/
-- 5a/
-- data2s.c
-- Assembler.nw
-- mkfile
-- builders/
-- Builder.nw
-- mk/
-- mkfile
-- compilers/
-- 5c/
-- 8c/
-- cc/
-- Compiler.nw
-- cpp/
-- mkfile
-- debuggers/
-- acid/
-- db/
-- Debugger.nw
-- libmach/
-- mkfile
-- Dockerfile
-- docs/
-- articles/
-- iw9p/
-- dosdisk.img
-- editors/
-- ed/
-- mkfile
-- include/
-- libc.h
-- ...
-- u.h
-- linkers/
-- 5l/
-- 8l/
-- Linker.nw
-- mkfile
-- readme.txt
-- windows/
-- libcomplete/
-- libplumb/
-- mkfile
-- rio/
-- Windows.nw

-- kernel/
-- arch/
-- concurrency/
-- console/
-- devices/
-- files/
-- filesystems/
-- init/
-- memory/
-- Kernel.nw
-- mkfile
-- network/
-- processes/
-- syscalls/
-- lib_core/
-- libbio/
-- libc/
-- Libcore.nw
-- libthread/
-- mkfile
-- lib_graphics/
-- Graphics.nw
-- libdraw/
-- libing/
-- libmemdraw/
-- libmemlayer/
-- mkfile
-- lib_networking/
-- lib9p/
-- libip/
-- mkfile
-- lib_strings/
-- libflate/
-- libregexp/
-- libstring/
-- mkfile
-- Makefile
-- mkconfig.pc
-- mkconfig.pi
-- mkfile
-- mkfile-host-Cygwin
-- mkfile-host-Linux
-- mkfile-host-macOS
-- mkfile-host-Plan9
-- MISC/
-- pc/
-- pi/

-- mkfiles/
-- 386/
-- arm/
-- makedirs
-- mkfile.proto
-- mklib
-- mkone
-- mkfile-target-pc
-- mkfile-target-pi
-- networking/
-- arp/
-- dhcp/
-- ftp/
-- http/
-- ip/
-- mkfile
-- ndb/
-- Network.nw
-- snoopy/
-- telnet/
-- profilers/
-- iostats/
-- mkfile
-- Profiler.nw
-- shells/
-- mkfile
-- rc/
-- Shell.nw
-- utilities/
-- archive/
-- byte/
-- calc/
-- compare/
-- files/
-- mkfile
-- pipe/
-- process/
-- text/
-- time/
-- Utilities.nw

-- ROOT/
-- 386/
-- arm/
-- lib/
-- rc/
-- tests/
-- usr/
    
```

Figure 1: <https://github.com/aryx/principia-softwarica> Layout.
(as generated by `tree -L 2 -F` and spread in 3 columns)

A New Way to Build

```

Goken Dockerfile excerpt
# Build goken on Ubuntu using
# gcc/binutils (and mk/rc)
FROM ubuntu:24.04

# Setup a basic C dev environment
RUN apt-get install -y gcc libc6-dev \
    byacc

# Now let's build from source
WORKDIR /src
COPY . .

# Small shell script (not GNU autoconf)
# to detect arch and generate
# ./mkconfig
RUN ./configure

# The script below obviously builds
# 'mk' (without needing mk) but also:
# - 'rc', which is called by 'mk'
# - 'ed', for the mkenam script
# run during the build
RUN ./scripts/build-mk.sh

# copy ./ROOT/<arch>/bin/{mk,rc,ed}
# to ./bin/
RUN ./scripts/promote-mk.sh

# make mk and rc accessible
ENV PATH="/src/bin:${PATH}"
ENV MKSHELL="/src/bin/rc"

# Let's build goken (using mk/rc built
# in the previous step)
RUN mk
RUN mk install

ENV PATH="/src/ROOT/amd64/bin:\
    /src/ROOT/arm64/bin:\
    ${PATH}"

Principia Softwarica Dockerfile
# Build principia on Ubuntu
# Linux (amd64 or arm64)
# for 386 (pc) and arm (pi)
FROM padator/goken

WORKDIR /principia
COPY . .

# 386
RUN cp mkconfig.pc mkconfig
RUN mk && mk install
RUN mk kernel


# arm
RUN cp mkconfig.pi mkconfig
RUN mk && mk install
RUN mk kernel

# VFAT tools
RUN apt-get install -y dosfstools \
    mtools

# making dosdisk.img
RUN dd if=/dev/zero of=tmp.img \
    bs=1M count=512
RUN mkfs.vfat tmp.img
RUN mcopy -i tmp.img -s -o \
    ROOT/* : :
RUN cat MISC/pc/bootsector tmp.img \
    > dosdisk.img
RUN rm -f tmp.img
    
```

Table 3: Principia Softwarica Build and Packaging Process.


Related Work



Project Oberon

Closest in spirit.


But Oberon-only, no networking.



Nand2Tetris

NAND gates to Tetris.

Principia: real OS, real code.



CS:APP

Explains concepts.

Principia: explains the source code.

Principia complements all three.

Future Work

- More explanations (LOE/LOC → 1)
AI could help write explanations
- Self-hosting (build Principia from Principia)
- Merge goken + principia codebases

Principia Softwarica

Principia removes the mental barrier.

The stack is not a black box.

These programs are not impossible to understand.

In other fields, students study the masters.

Read the masters' code. Understand everything.

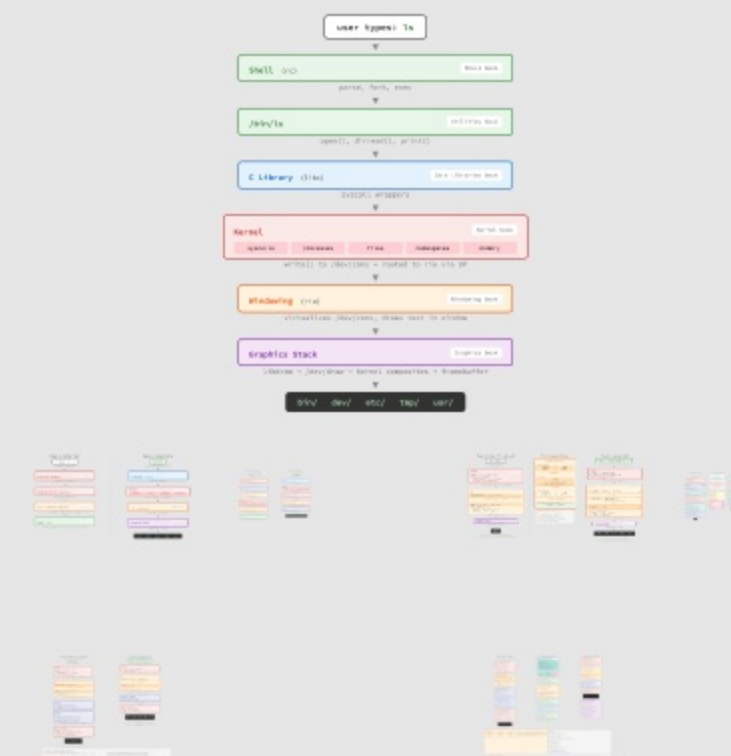
Become a better engineer.

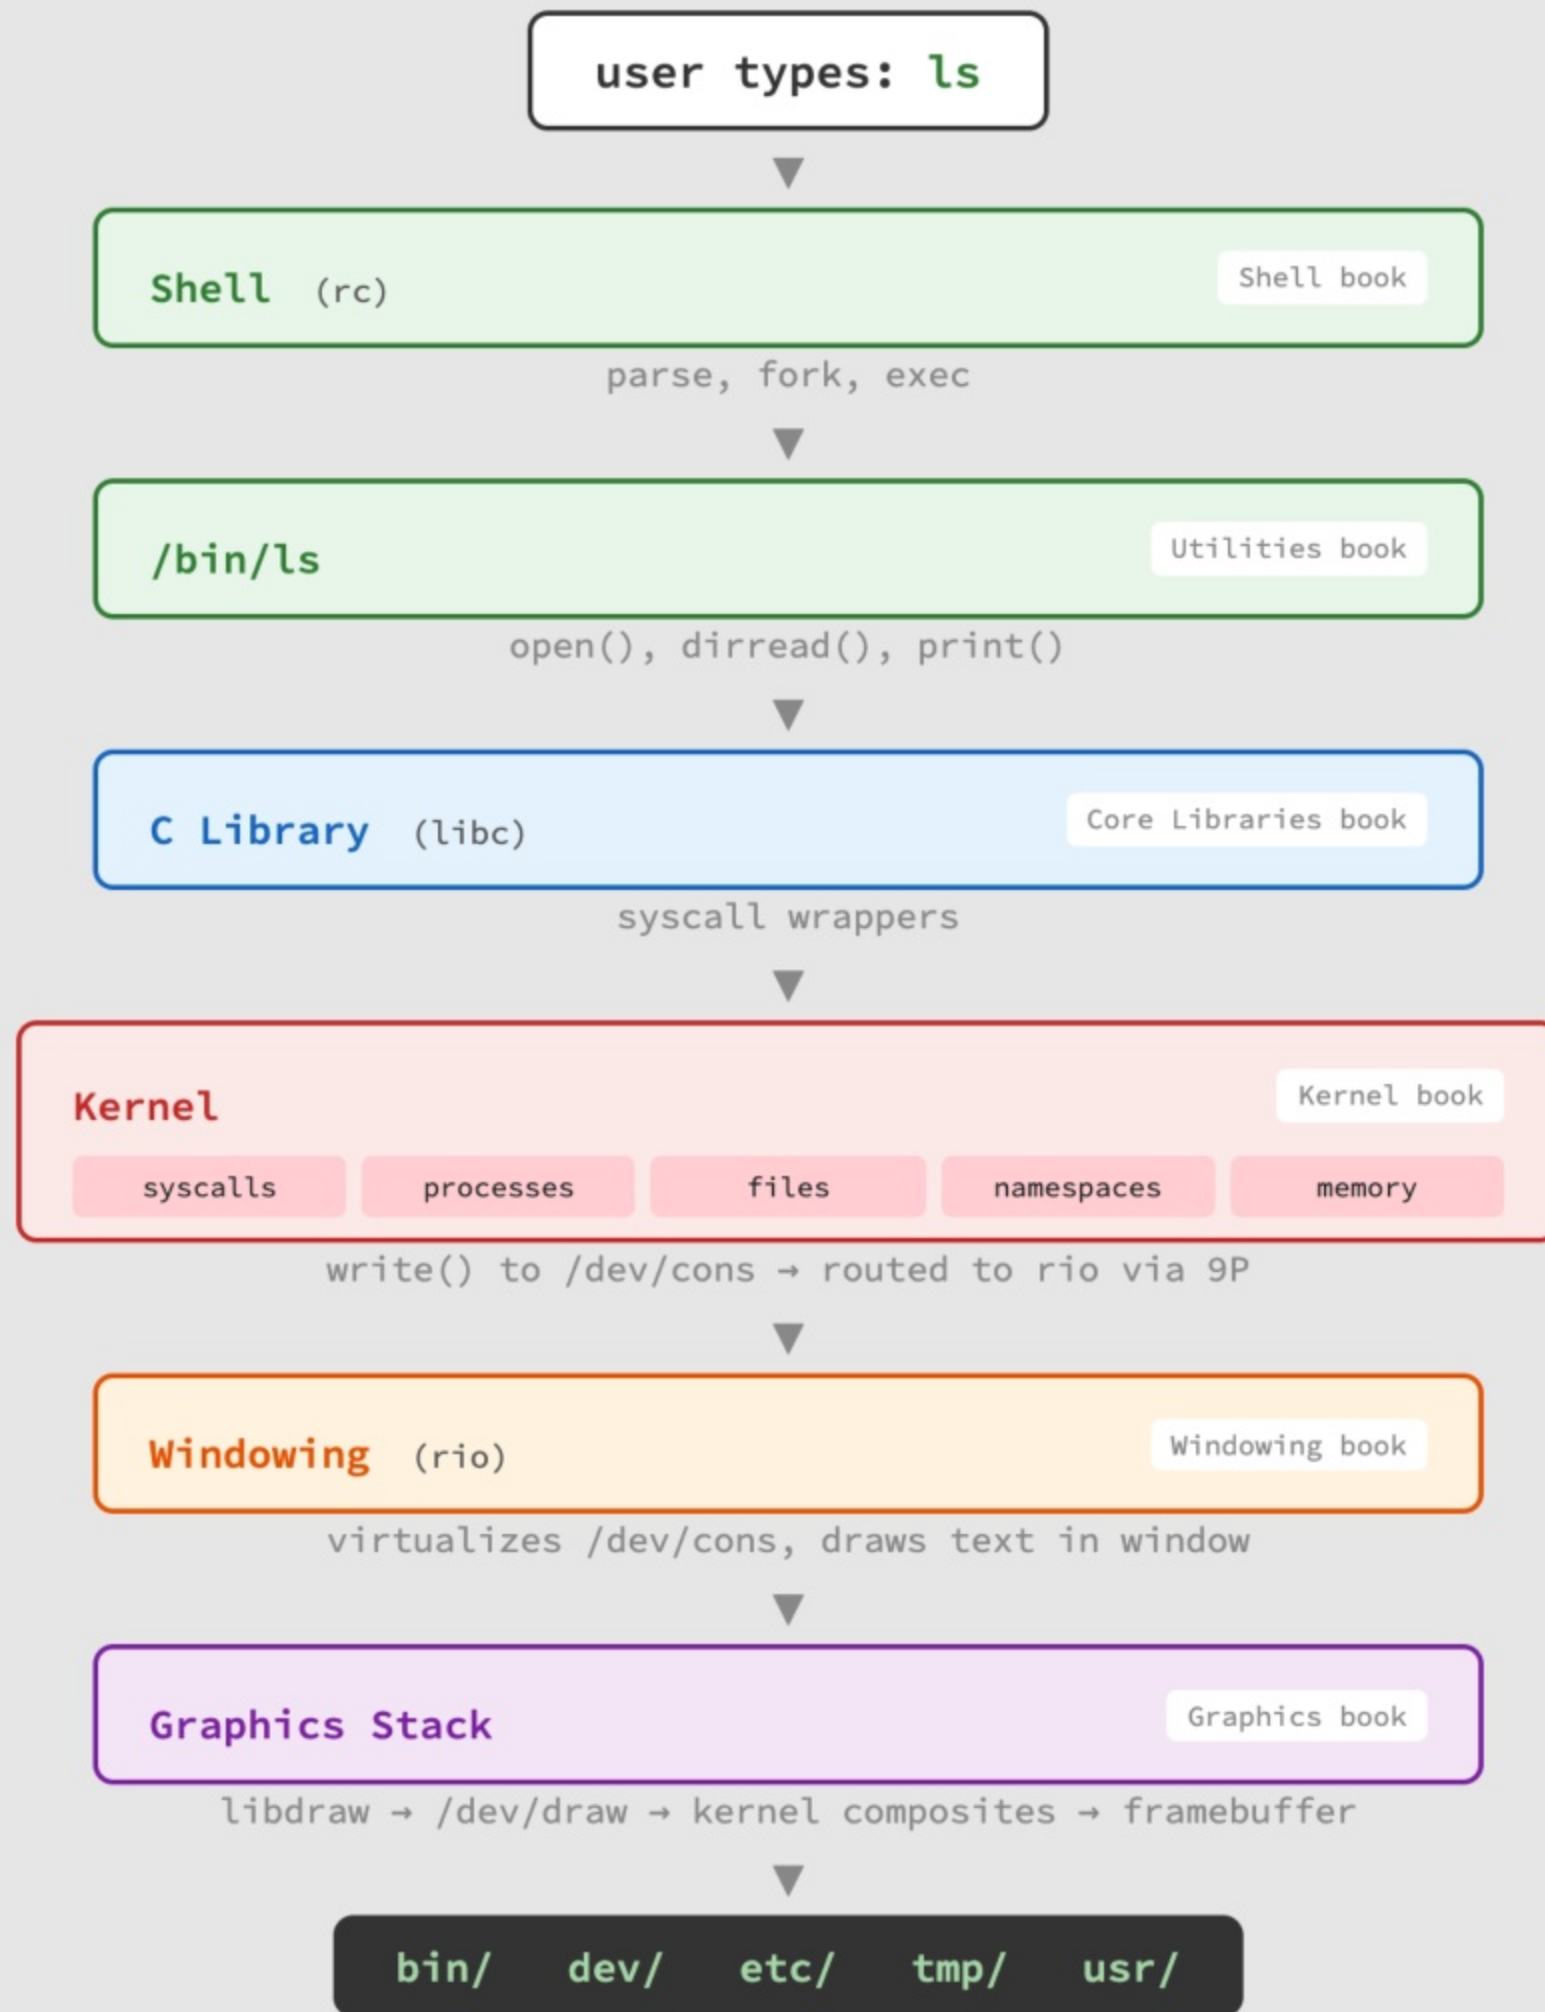
"What happens when you type ls?"

After reading Principia: you'll know. Fully.

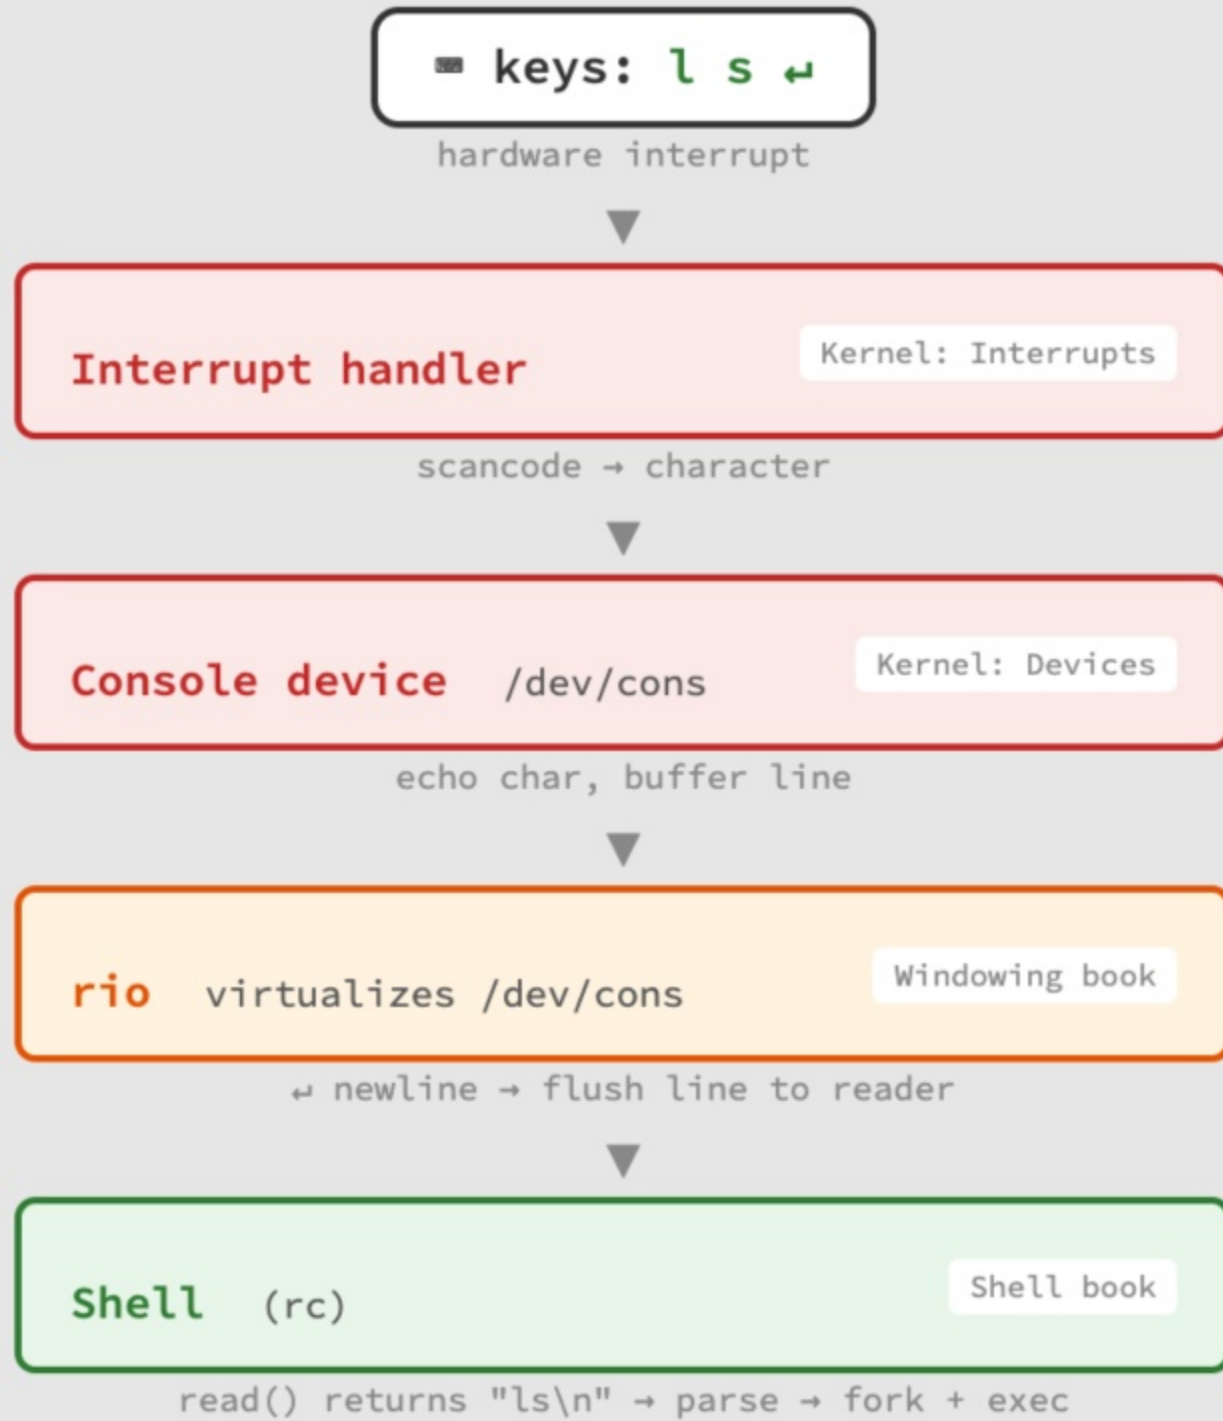
Clone it. Build it. Read the books.

github.com/aryx/principia-softwarica

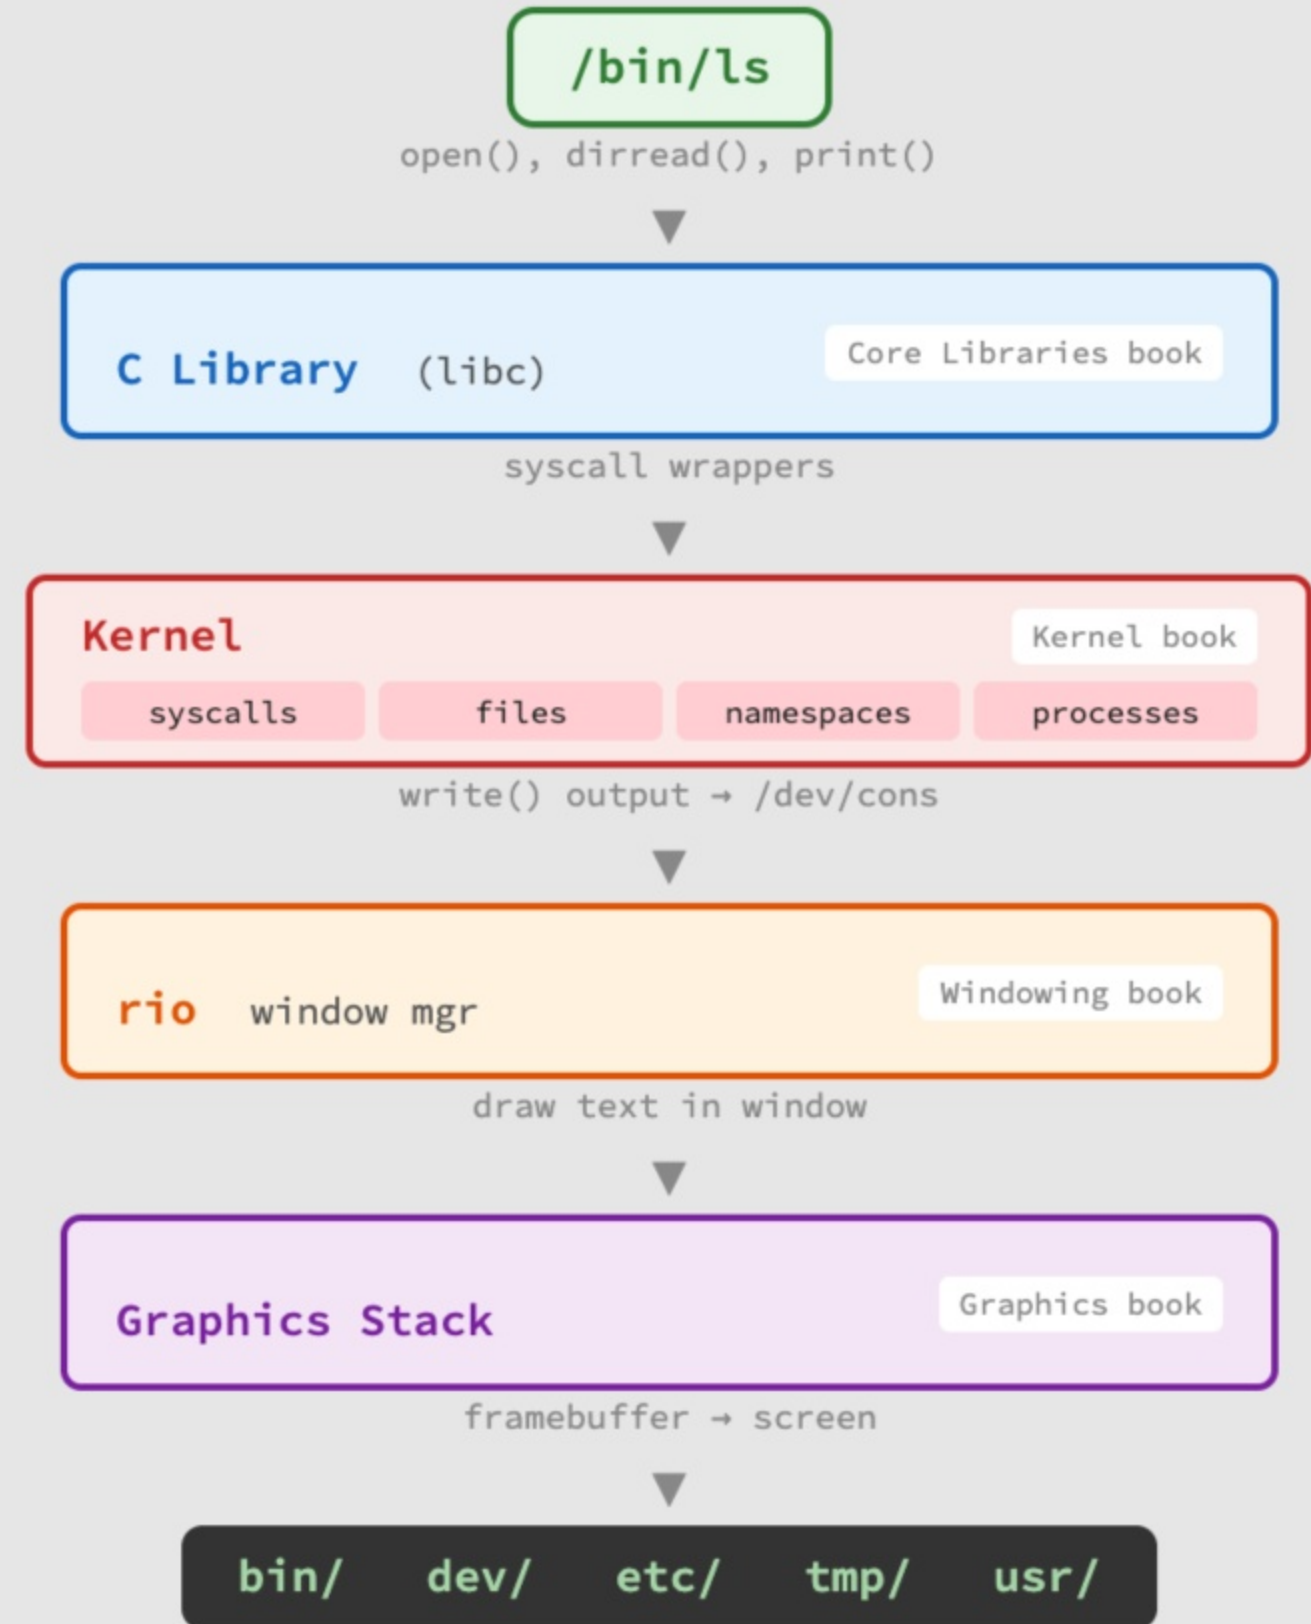




Phase 1: Typing "ls"



Phase 2: Executing ls



Phase 1: Typing "ls"

keys: l s ↵

USB/PS2 interrupt

Input subsystem evdev

Kernel

scancode → keycode

Xorg X server

~500K LOC

KeyPress event → xterm

xterm terminal emulator

88K LOC

write 'l' to PTY master fd

Kernel: PTY + TTY

Kernel

pty driver

n_tty ldisc

echo

echo back to master + buffer line

xterm VT100 parse + render

88K LOC

Xft/FreeType → X11 draw request

bash (shell)

180K LOC

read() returns "ls\n" → parse → fork + exec

Phase 2: Executing ls

/bin/ls

opendir(), readdir(), printf()

glibc (C library)

1.5M LOC

syscall wrappers

Kernel

~40M LOC

syscalls

VFS

ext4/btrfs

scheduler

write() output → PTY slave fd

PTY + TTY n_tty ldisc

Kernel

data → master side

xterm VT100 + render

88K LOC

Xft/FreeType → X11 draw request

Xorg composite + render

~500K LOC

DRM/KMS → framebuffer

bin/ dev/ etc/ tmp/ usr/

Phase 1: Typing "l" – echo path

key press: `l`

hardware interrupt

Kernel

Kernel book

`kbdputc('l') → circular buffer`
`clock tick → echo():`
`qproduce(kbdq, 'l') – into queue`
raw mode → skip screen echo

`read /dev/cons returns 'l' from kbdq`

rio

Windowing book

`keyboardthread – recv(keyboardctl→c)`

`sendp(input→ck, 'l') channel`

`winctl (per window) – event loop`

`WKey → wkeyctl(w, 'l')`
`→ winsert() – insert rune in w→r[]`
`→ frinsert() – draw 'l' glyph`

`draw() via /dev/draw`

Graphics Stack

Graphics book

`framebuffer → screen`

`% l`

Then 's' and '␣' follow the same path.
On '␣', `winctl` marks line ready for shell to read.

Rio's internal threads

Thread architecture

Global threads:

keyboard thread

mouse thread

↓ channels ↓

Per window:

winctl

event loop: `alt(WKey, WMouse, WCreed, WCwrite, WCtl, ...)`

↑ channels

9P file server:

filsysproc

reads 9P msgs

xfidctl

worker pool

⚡ 9P over /dev/cons

Per window (separate process):

`winshell → rc (shell)`

Channels (typed message queues):

`ck` – keyboard runes → `winctl`
`mc.c` – mouse events → `winctl`
`consread` – shell read req ↔ `winctl`
`conswrite` – shell write req ↔ `winctl`
`cctl` – control msgs → `winctl`

Phase 2: ls output path

`/bin/ls → print("bin/\n...")`

`write() syscall to virtual /dev/cons`

Kernel

Kernel book

`sys_write() → devtab→write()`
but `/dev/cons` is rio's virtual file
→ 9P write message to rio

9P write to rio's file server

rio

Windowing book

`filsysproc – read9pmsg() → dispatch`

`xfidctl worker → xfidwrite(Qcons)`

send on `conswrite` channel

`winctl (per window)`

`WCwrite → recv runes`
→ `winsert()` – insert into `w→r[]`
→ `frinsert()` – draw text

`draw() via /dev/draw`

Graphics Stack

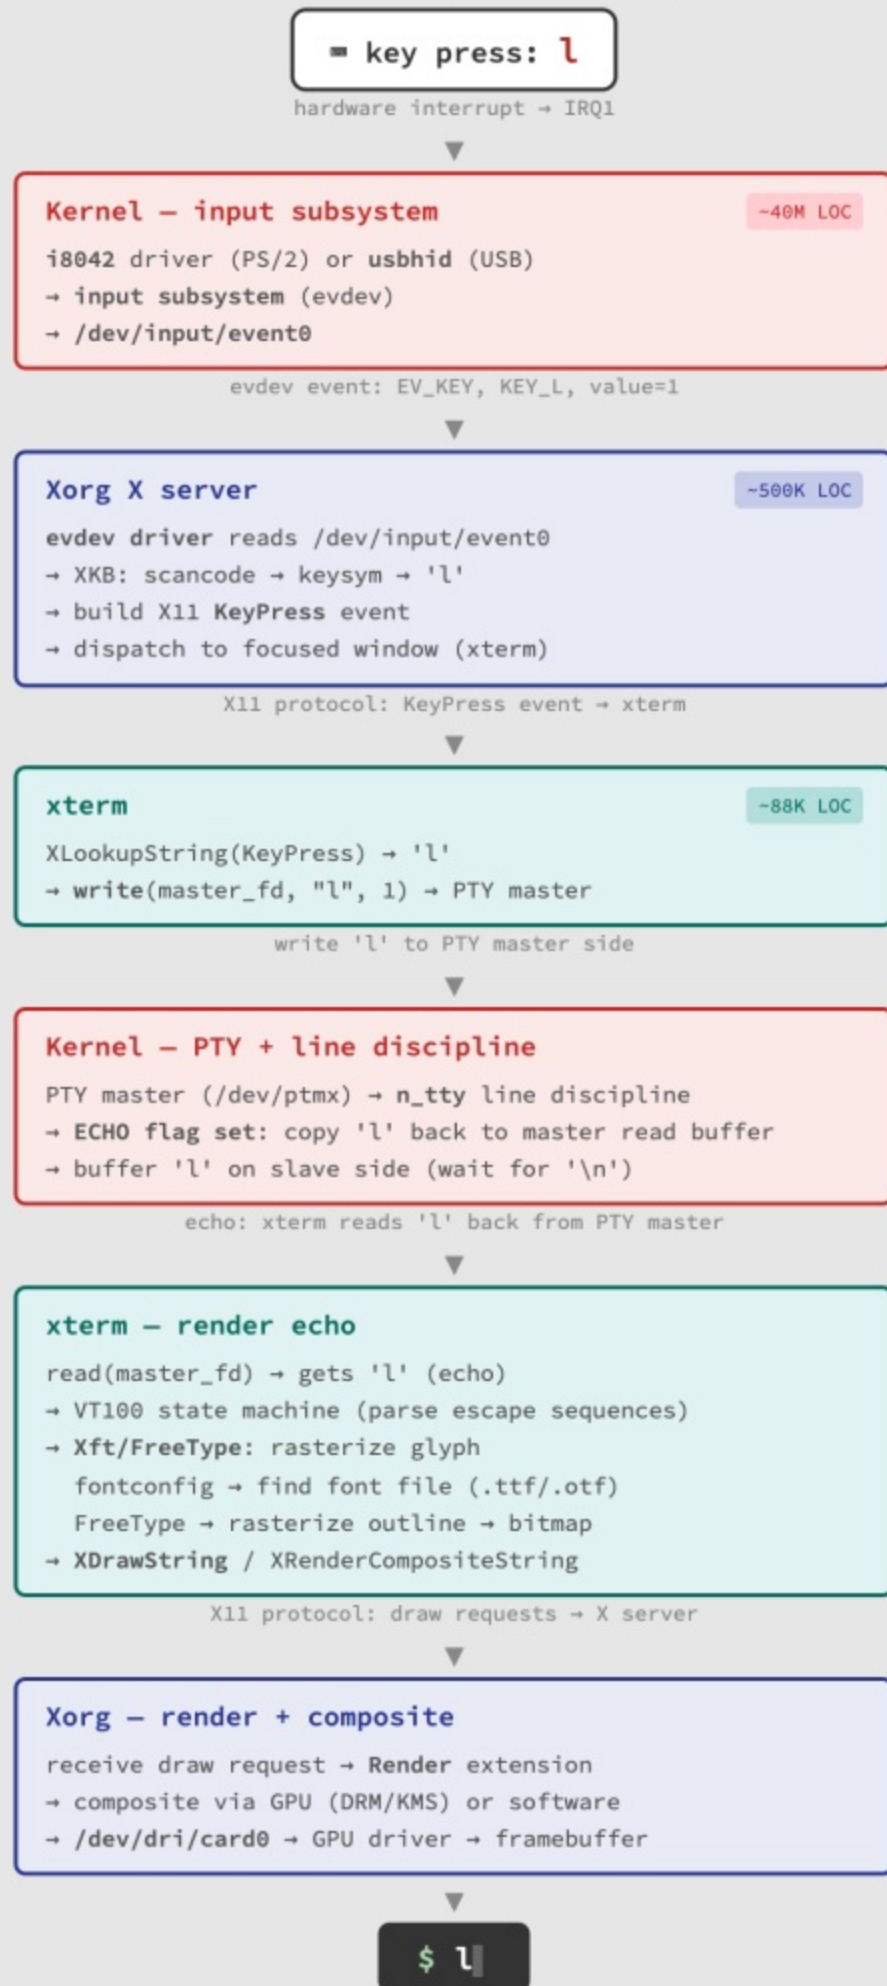
Graphics book

`framebuffer → screen`

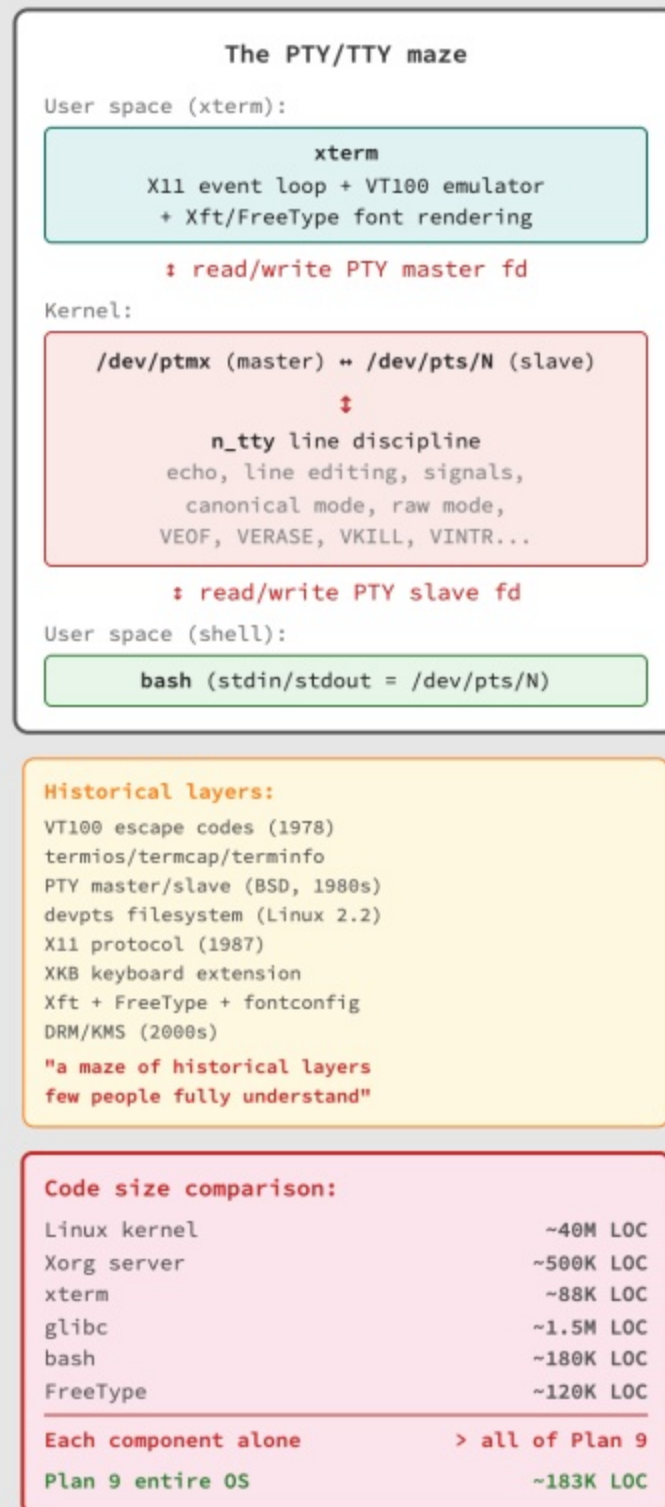
`bin/ dev/ etc/ tmp/ usr/`

Both echo and output go through
the same `winctl → winsert → frinsert` path.

Phase 1: Typing "l" – echo path



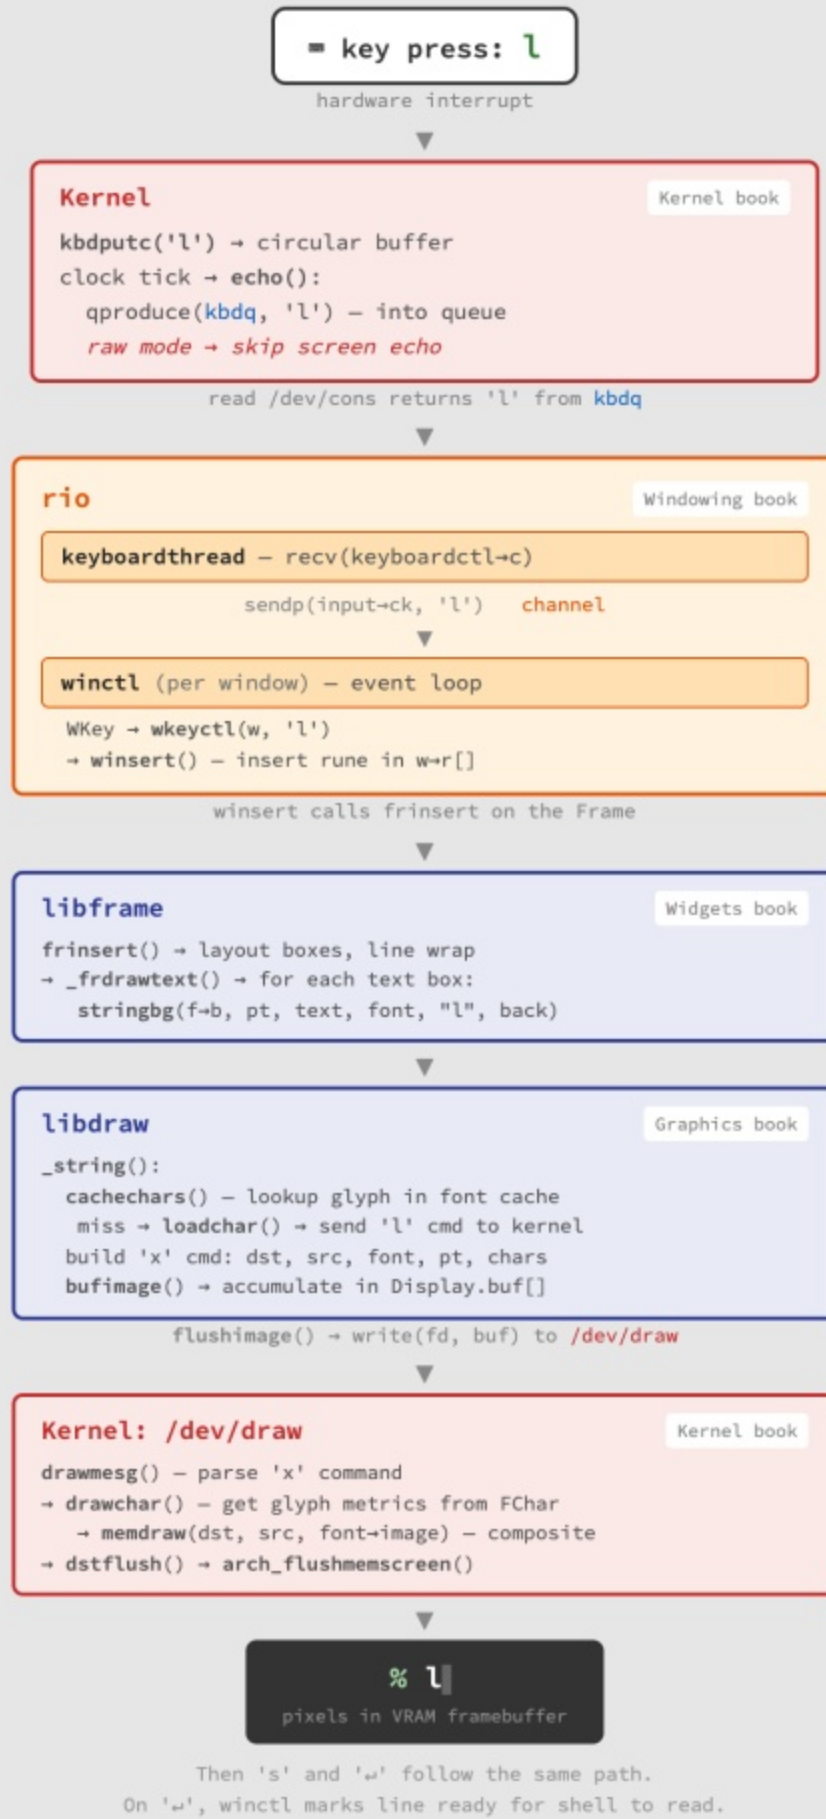
Linux terminal architecture



Phase 2: ls output path



Phase 1: Typing "l" – echo path



Phase 2: ls output path



Key insight: rio goes through the kernel twice

1st pass (input): keyboard interrupt → kbdq → /dev/cons → rio

2nd pass (display): rio → libdraw protocol → /dev/draw → kernel composites → framebuffer

Font system: cachechars() keeps glyph cache in user-space (libdraw).
On cache miss, loadchar() sends subfont pixels to kernel via /dev/draw 'l' cmd.
Kernel stores glyph bitmaps in font-image (Memimage in kernel memory).

PHASE 1: TYPING "l" - ECHO PATH

= key press: **l**

PS/2 scancode → IRQ: hardware interrupt

Kernel - interrupt handler + console device Kernel book

18042intr() = scancode to runs
 → kbdputc('l') = kbd.tstbuf[] circular buffer
 → batched at interrupt time, no queue ops yet → clock tick (22ns) = kbdputclock():
 = echo('l');
 qproduce(kbdq, 'l') → into kernel queue
 kbd.raw == true = SKIP echoscreen()
 rto set raw mode = kernel does NOT echo

110thread keyboard reader: read(/dev/console) = gets 'l' from kbdq

rio - keyboard input to window Windowing book

keyboardthread = recv(keyboardctl+q, 'l')

winctl (per window) = event loop: alt(WKey, WMouse, WCreat, ...)

case WKey:
 = wkeyctl(w, 'l') = not raw, not navigation key
 = winert(w, &l, l, wqh) = insert rune in w[] buffer
 runemove(w+@, &r, l) = shift text, insert
 adjust cursors: w=qh+, w=ql+

winsert calls frinsert on the visible frame

libframe - text widget rendering Widgets book

frinsert(dw+fr, "l", qh = w+rg)
 → bxcas() → break text into frbox[] (text runs = tabs + newlines)
 → calculate layout, line wrapping
 → _fdrawtext(f, pt, text_color, back_color):
 for each box: string(f+@, pt, text, ZF, f-font, "l", back, ZF)

libdraw - draw protocol + font cache Graphics book

string@ = _string(dw, pt, arc, font, "l", bg, SoverD);
 cachechars(font, "l") = hash lookup in font.cache[]
 HIT = use cached index
 miss = loadchar() = send 'l' cmd to kernel
 (load subfont glyph bitmap into font-cache+img)
 build 'x' cmd: dst_id, src_id, font_id, pt, clipr, char_indices
 bufimage(display, m) = append to Display_buf[32KB]

flushimg() = write(fd, buf) to /dev/draw

Kernel - /dev/draw graphics device Kernel book

devdraw write() = drawmsg() = parse binary 'x' command
 unpack: dst, src, font, pt, clipr, char_indices
 fill background: memdraw(dst, r, bg, ...)
 for each char index:
 = drawchar(dst, bg, q, src, font, cl, SoverD)
 fc = font-fchar[cl] = glyph metrics (left, width, winy, hxy)
 = memdraw(dst, r, src, op, font+img, Pt(fc+wfx, fc+winy), SoverD)
 glyph bitmap is the MASK = src color painted through it
 = dstflush(dw_id, dst, r) = arch_flushmemscreen(r)



Some path repeats for 's' and 'u'. After Enter: winctl has "ls\n" in w[]. Line ready for reader.

THE BRIDGE: SHELL READS "ls\n"

Namespace setup (at window creation) Kernel book

winshell() → rio creates a process for each window:
 1. rfork(RFNAMEG | RFFDG | RFDW) → new per-process namespace (Pgrp)
 = newproc() = copy parent's mount table
 = pgrpncpy() = copy parent's mount table
 2. Pflaymount(Pflays, w+id) = mount(fd, -1, "/mnt/ways", MNTPL, 0)
 0P connection to rio's file server = /mnt/ways/
 = bind("/mnt/ways", "/dev", MNTFSMS)
 union mount: rio's files searched BEFORE real /dev
 3. close(0); open("/dev/console", OREAD) = stdin
 close(1); open("/dev/console", OWRITE) = stdout
 /dev/console now resolves to rio's Qcons, not kernel's!
 4. execve("/bin/rsh", "/bin/rsh", args) = launch shell

Shell's namespace (per-process)



Now the shell is blocked on read("/dev/console")...

Shell (rc) - blocked read Shell book

read(0, buf, 4) = stdin = rio's /dev/console
 = libc:_read() = SYSCALL = kernel sys_read()

kernel routes to rio via 0P: Tread = rio's pipe

rio - 0P file server handles read Windowing book

fllayproc = readmsg() = Tread for Qcons

xfidctl worker = xfidread(Qcons)

send on console channel = blocks until winctl ready

winctl = case WCreat (after 'u' made line ready)

recv(crm.c, &ptr) = get buffer from xfidread
 copy from w=[wqh..] until '\n':
 runetochar(l, &w=[wqh+]) = "l", "s", "u"
 send(crm.c, &ptr) = return "ls\n" to xfidread

0P Tread = kernel = shell's read() returns 3 bytes: "ls\n"

Shell (rc) - parse + execute Shell book

yyparse() = parse "ls\n" → AST: simple command
 rfrk(RFFRDC|RFFDG|RFDTEG) = fork child process
 child inherits namespace = same /dev/console = same rio window
 exec("/bin/ls", ["ls"]) = replace child with ls binary

libc: exec() = SYSCALL = kernel sys_exec()

C Library (libc) Core Libraries book

ls binary linked with libc: open(), dirread(), printf()
 Each wraps a syscall: libc = SYSCALL instruction = kernel trap

syscall wrappers = kernel

/bin/ls - the program Utilities book

fd = open("/", OREAD) = kernel: namec("/") = root devtab
 dirread(fd, dir) = kernel: devtab-read[] = directory entries
 for each entry: printf("%s\n", name)
 = write(1, "bin/\n", 4) = stdout = rio's /dev/console

write() to /dev/console = enters the OUTPUT path =

PHASE 2: ls OUTPUT → SCREEN

Kernel - syscall + 0P routing Kernel book

sys_write(fd=1, "bin/\n", 4)
 = fd 1 = Chan = mounted filesystem check
 = /dev/console is rio's Qcons (via namespace bind)
 = devant: build 0P write message
 = write to rio's 0P pipe

0P Twrite = rio's file server pipe

rio - 0P file server handles write Windowing book

fllayproc = readmsg() = Twrite for Qcons

xfidctl worker = xfidwrite(Qcons)

cuttunes("bin/\n", r) = bytes = Rune[]

send on console channel = runs to winctl

winctl = case WWrite

recv(crm.cw, &ptr) = get runes
 handle 'b' backspace processing
 = winsert(r, q, nr, wqh) = insert output text
 w=qh == nr = advance read head past output

winsert calls frinsert on the frame

libframe + libdraw - same path as echo Graphics + Widgets

frinsert() = _fdrawtext() = string@()
 = _string() = cachechars() = build 'x' command
 = bufimage() = flushimg() = write to /dev/draw

Kernel - /dev/draw

drawmsg() = drawchar() = memdraw()
 = dstflush() = arch_flushmemscreen(r)

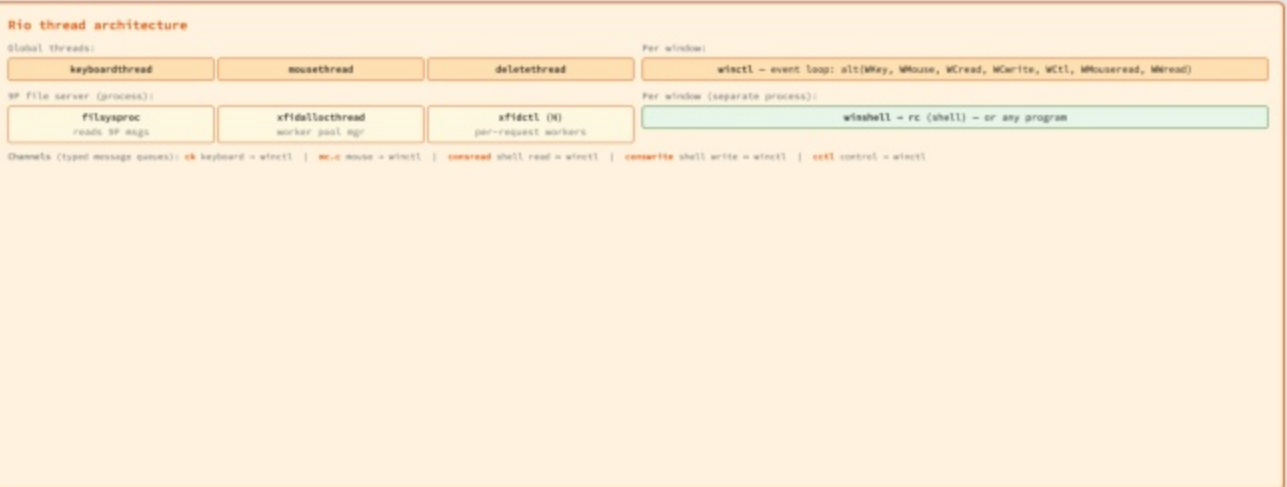
Font system - how glyph bitmaps work

Font files: /lib/fonts/bits/unicode/0-font
 meta-file listing subfont ranges:
 0x0000-0x00FF Latin-0 = covers ASCII
 0x0100-0x02AF Latin-1 = extended Latin

subfont: Latin-0 = bitmap + metrics
 image: all glyphs side by side in one image
 Fontchar[]: (x, top, bottom, left, width) per glyph

Cache: font-cache+img (user-space image)
 cachechars('l') = hash lookup = cache index
 miss = loadchar(): read subfont, send glyph bitmap to kernel via 'l' cmd = kernel stores in font+img

Rendering: drawchar(dst, src, font+img)
 memdraw(dst, r, src, op, mask+font+img, op, SoverD)
 glyph bitmap = alpha mask
 pixel = src_color * mask + dst_color * (1 - mask)
 1-bit font: binary mask = fast chardraw()
 anti-aliased: gray mask = alpha+draw()



Key insights

Rio goes through the kernel twice:
 1st: keyboard IRQ = kbdq = /dev/console = rio
 2nd: rio = /dev/draw = kernel composites = framebuffer

Namespaces are the key abstraction:
 Shell's /dev/console = kernel's /dev/console
 Per-process mount table (Pgrp) = bind("/mnt/ways", "/dev", MNTFSMS)
 Same mechanism that Docker/containers use (Linux namespaces = Plan 9 idea)

Everything is a file + 0P:
 Console I/O = read/write on /dev/console
 Graphics = read/write on /dev/draw
 Mouse = read on /dev/mouse
 All go through the same syscall path

winctl is the heart of rio:
 Single event loop per window (like a mini server)
 Both echo (WKey) and output (WWrite) = same winsert = frinsert path

User-space builds commands, kernel renders pixels:
 libdraw accumulates draw protocol to 32KB buffer
 kernel /dev/draw does actual pixel compositing in memdraw()
 Font glyphs used as alpha masks = elegant Porter-Duff compositing

Principia Softwarica

Plan 9 Code Explained

Yoann Padioleau

IWP9



The Best Engineers Understand the Stack

Facebook linker slow, but benchmark for everyone
 → most engineers: stack, accepted it, black box
 We had to wait for [Ian Lance Taylor](#) (Google, 2006) to write `gold` → a new linker for GNU Binutils.
 Not a young hotshot. Old school Unix/GCC hacker (since 1996).
 Also part of the Go team at Google.
 The kind of engineer who understands the full stack.
 Understand deeply = build better

The Education Gap

"What happens when you type `ls` in a terminal window?"
 Keith Adams → colleague at Facebook, later Chief Architect at Slack → used this as his interview question.
 Simple question. The answer involves: the shell, the C library, the kernel, the graphics stack, the windowing system.
 Most engineers can't answer.
 "Full stack" today = React + Node + cloud.
 REAL full stack = compiler, linker, kernel, ops/etc

The AI Era Makes This More Important

We write less code. We need MORE understanding.
 AI generates code → but who evaluates it?
 Who debugs the segfault? The linker error?
 AI tools like `CLAUDE` Code USE the classic commands:
`grep`, `ls`, `diff`, `awk`, `gcc`, `ls` ...
 The AI is a power user of the very programs Principia explains.
 Writing code = 20%. Understanding = 80%.
 AI handles the 20%. The 80% becomes everything.

But the Stack is Enormous

Linux kernel -40M LOC
 gcc -15M LOC
 glibc -1.5M LOC
 bash -1.6M LOC
 Borg server -50M LOC
 stars above -8M LOC

Each component: more code than ALL of Plan 9

Plan 9 - 100k Lines Code

Principia Software
 Principia is a software stack for building systems.
 It is a collection of tools and libraries that work together to create a complete operating system.
 Principia is designed to be easy to use and to integrate with existing systems.
 Principia is a powerful user of the very programs Principia explains.
 Writing code = 20%. Understanding = 80%.
 AI handles the 20%. The 80% becomes everything.

A Series of Books

| Category | Book | Program(s) | LOC | LOE | Pages |
|-----------------------|------------------|---|---------------|--------------|-------------|
| Core system | Kernel | 9pi | 35000 | 3200 | 732 |
| | Core libraries | libc libregexp libthread | 19000 | 1600 | 438 |
| | Shell | rc | 6500 | 1700 | 166 |
| Development toolchain | C compiler | 5c libcc | 18500 | 1900 | 471 |
| | Assembler | 5a | 3600 | 4400 | 176 |
| | Linker | 5l | 7500 | 5400 | 296 |
| Developer tools | Editor | ed | 1600 | 200 | 45 |
| | Build system | mk | 4350 | 4050 | 197 |
| | Debuggers | db acid strace libmach | 13100 | 1000 | 321 |
| Graphics | Graphics stack | prof time kprof stats iostats | 3900 | 350 | 102 |
| | | /dev/draw | 18500 | 3400 | 507 |
| | | libdraw libmemdraw libmemlayer | | | |
| Networking | Windowing system | rio libframe libcomplete libplumb | 8800 | 4000 | 289 |
| | | /dev/net | 18300 | 4800 | 457 |
| | | libip lib9p | | | |
| Misc | CLI utilities | cat ls grep sed diff tar gzip ... | 23900 | 650 | 493 |
| Total: | | | 182550 | 36650 | 4690 |

Table 1: Principia Softwarica Books and Their Statistics.
 (LOC = lines of code; LOE = lines of explanation; Pages = number of typeset pages)

A New Plan 9 Fork

```

Why a New Fork?
https://github.com/aryx/principia-softwarica

Repository structure
Screenshot: Repository structure diagram showing folders like assemblers/, kernel/, mkfiles/, builders/, compilers/, debuggers/, docs/, editors/, linkers/, windows/, include/, windows/.

Screenshot: Repository structure diagram showing folders like assemblers/, kernel/, mkfiles/, builders/, compilers/, debuggers/, docs/, editors/, linkers/, windows/, include/, windows/.

Screenshot: Repository structure diagram showing folders like assemblers/, kernel/, mkfiles/, builders/, compilers/, debuggers/, docs/, editors/, linkers/, windows/, include/, windows/.

-- assemblers/
|-- 5a/
|-- 5a/
|-- data2s.c
|-- Assembler.nw
|-- mkfile
-- builders/
|-- Builder.nw
|-- mk/
|-- mkfile
-- compilers/
|-- 5c/
|-- 5c/
|-- cc/
|-- Compiler.nw
|-- cpp/
|-- mkfile
-- debuggers/
|-- acid/
|-- db/
|-- Debugger.nw
|-- libmach/
|-- mkfile
-- Dockerfile
-- docs/
|-- articles/
|-- iw9p/
-- editors/
|-- ed/
|-- mkfile
-- include/
|-- libc.h
|-- ...
|-- u.h
-- linkers/
|-- 5l/
|-- 5l/
|-- Linker.nw
|-- mkfile
-- readme.txt
-- windows/
|-- libcomplete/
|-- libplumb/
|-- mkfile
|-- rio/
|-- Windows.nw

-- kernel/
|-- arch/
|-- concurrency/
|-- console/
|-- devices/
|-- files/
|-- filesystems/
|-- init/
|-- memory/
|-- Kernel.nw
|-- mkfile
-- network/
|-- processes/
|-- syscalls/
-- lib_core/
|-- libbio/
|-- libc/
-- libcore.nw
|-- libthread/
|-- mkfile
-- lib_graphics/
|-- Graphics.nw
|-- libdraw/
|-- libing/
|-- libmemdraw/
|-- libmemlayer/
|-- mkfile
-- lib_networking/
|-- lib9p/
|-- libip/
|-- mkfile
-- lib_strings/
|-- libflate/
|-- libregexp/
|-- libstring/
|-- mkfile
-- Makefile
|-- mkconfig.pc
|-- mkconfig.pi
|-- mkfile
-- mkfile-host-Cygwin
-- mkfile-host-Linux
-- mkfile-host-macOS
-- mkfile-host-Plan9
-- MISC/
|-- pc/
|-- pi/

-- mkfiles/
|-- 386/
|-- arm/
|-- makedirs
|-- mkfile.proto
|-- mklib
|-- mkone
-- mkfile-target-pc
-- mkfile-target-pi
-- networking/
|-- arp/
|-- dhcp/
|-- ftp/
|-- http/
|-- ip/
|-- mkfile
|-- ndb/
-- Network.nw
|-- snoopy/
|-- telnet/
-- profilers/
|-- iostats/
|-- mkfile
|-- Profiler.nw
-- shells/
|-- mkfile
|-- rc/
|-- Shell.nw
-- utilities/
|-- archive/
|-- byte/
|-- calc/
|-- compare/
|-- files/
|-- mkfile
|-- pipe/
|-- process/
|-- text/
|-- time/
|-- Utilities.nw
-- ROOT/
|-- 386/
|-- arm/
|-- lib/
|-- rc/
|-- tests/
|-- usr/
    
```

Figure 1: <https://github.com/aryx/principia-softwarica> Layout.
 (as generated by `tree -L 2 -F` and spread in 3 columns)

A New Way to Build

```

Goken Dockerfile excerpt
# Build goken on Ubuntu using
# gcc/binutils (and mk/rc)
FROM ubuntu:24.04

# Setup a basic C dev environment
RUN apt-get install -y gcc libc6-dev \
    byacc

# Now let's build from source
WORKDIR /src
COPY . .

# Small shell script (not GNU autoconf)
# to detect arch and generate
# ./mkconfig
RUN ./configure

# The script below obviously builds
# 'mk' (without needing mk) but also:
# - 'rc', which is called by 'mk'
# - 'ed', for the mkenam script
# run during the build
RUN ./scripts/build-mk.sh

# copy ./ROOT/<arch>/bin/{mk,rc,ed}
# to ./bin/
RUN ./scripts/promote-mk.sh

# make mk and rc accessible
ENV PATH="/src/bin:${PATH}"
ENV MKSHELL="/src/bin/rc"

# Let's build goken (using mk/rc built
# in the previous step)
RUN mk
RUN mk install

ENV PATH="/src/ROOT/amd64/bin:\
    /src/ROOT/arm64/bin:\
    ${PATH}"

Principia Softwarica Dockerfile
# Build principia on Ubuntu
# Linux (amd64 or arm64)
# for 386 (pc) and arm (pi)
FROM padator/goken

WORKDIR /principia
COPY . .

# 386
RUN cp mkconfig.pc mkconfig
RUN mk && mk install
RUN mk kernel

# arm
RUN cp mkconfig.pi mkconfig
RUN mk && mk install
RUN mk kernel

# VFAT tools
RUN apt-get install -y dosfstools \
    mtools

# making dosdisk.img
RUN dd if=/dev/zero of=tmp.img \
    bs=1M count=512
RUN mks.vfat tmp.img
RUN mcopy -i tmp.img -s -o \
    ROOT/* : :
RUN cat MISC/pc/bootsector tmp.img \
    > dosdisk.img
RUN rm -f tmp.img
    
```

Table 3: Principia Softwarica Build and Packaging Process.

Take this with you. Revisit anytime.

Missed something? Want to explore further?
Scan or click below to open this presentation.
Anytime, anywhere.

[View presentation](#)

