

# 9 FRONT KÄPTN BLAUBÄR

4661



This book was typeset (`troff -ms -mpictures |lp -dstout |ps2pdf`) in Lucida Sans by the author, using a Lenovo ThinkPad X60 Tablet running the 9front operating system.

Rendered: 2015-7-19

9FRONT  
9front.org

This is a work of fiction. Names, characters, places and incidents either are the product of the author's imagination or are used fictitiously, and any resemblance to actual persons, living or dead, businesses, companies, events or locales is entirely coincidental.

MIT/CC0/Public Domain

## 9FRONT FREQUENTLY QUESTIONED ANSWERS

# 9front System

*Those who can do, those who can't write and those who can't write make ezines.*  
— Sape Mullender

**ACHTUNG!** This document's STRUCTURE — and some of its TEXT — is plagiarized! — from the OPENBSD FAQ!!

**ACHTUNG!** Information provided by this document may be outdated or just plain wrong. Use your brain. NO REFUNDS.



## **0 – Introduction to Plan 9**

0.1 – What is Plan 9?

0.1.1 – Plan 9 is not UNIX

0.1.1.1 – Plan 9 is not plan9port

0.1.1.2 – Plan 9 is not Inferno

0.1.2 – Plan 9 is not a product

0.1.3 – Plan 9 is not for you

0.2 – Why Plan 9?

0.2.1 – What do people like about Plan 9?

0.2.1.1 – What do you use Plan 9 for?

0.2.2 – What do people hate about Plan 9?

0.2.2.1 – What is not in Plan 9

0.2.3 – Why did Plan 9's creators give up on Plan 9?

0.2.4 – What is the deal with Plan 9's weird license?

0.2.4.1 – Richard Stallman hates the Plan Nine license (circa 2000)

0.2.4.2 – Theo de Raadt hates the Plan 9 license (circa 2003)

0.2.4.3 – Everyone hates the Plan 9 license (circa 2014)

0.3 – Further Reading

0.3.1 – Plan 9 papers

0.3.2 – Man pages

0.3.3 – Web pages

0.3.4 – Books

## **1 – Introduction to 9front**

1.1 – What is 9front?

1.1.1 – Cirno's been our mascot for a while.

1.2 – On what systems does 9front run?

1.3 – Why might I want to use 9front?

1.3.1 – New Features

1.3.1.1 – New Programs

1.3.1.2 – New Hardware Support

1.4 – Is 9front really free?

1.5 – How can I help support 9front?

1.6 – Who maintains 9front?

1.7 – When is the next release of 9front?

1.8 – What is included with 9front?

1.9 – Can I use 9front as a desktop system?

1.10 – Why is/isn't ProductX included?

1.11 – Fine, where can I get 9front?

1.11.1 – Mirrors

1.11.1.1 – 9front.iso

## **2 – Getting to know 9front**

2.1 – Web Pages

2.2 – Mailing Lists

2.2.1 – Mailing List Archives

2.2.2 – Other useful mailing lists

2.3 – Manual Pages

2.3.2 – How do I write my own manual page?

2.4 – Reporting Bugs

2.4.1 – Your bug report SUCKS

2.4.2 – How do I get more useful info for developers?

2.4.2.1 – Take a photo of the panic screen

2.4.2.2 – Make a stack trace of a crashed process

### **3 – Hardware**

#### 3.1 – Selecting Hardware

#### 3.2 – Known Working Hardware

##### 3.2.1 – Audio

##### 3.2.2 – Graphics

###### 3.2.2.1 – AGP

###### 3.2.2.2 – Integrated

###### 3.2.2.3 – PCI Express

##### 3.2.3 – Networking

###### 3.2.3.1 – Ethernet

###### 3.2.3.1.1 – Integrated

###### 3.2.3.1.2 – USB

###### 3.2.3.2 – WiFi

###### 3.2.3.2.1 – Bridge (external)

###### 3.2.3.2.2 – Mini-PCI

###### 3.2.3.2.3 – Mini-PCI Express

###### 3.2.3.2.4 – PCI

###### 3.2.3.2.5 – PCMCIA

##### 3.2.4 – Tablet Digitizers

###### 3.2.4.1 – Serial

###### 3.2.4.1.1 – Integrated

##### 3.2.5 – Laptops

###### 3.2.5.1 – Acer

###### 3.2.5.2 – IBM/Lenovo

###### 3.2.5.2.1 – ThinkPad

###### 3.2.5.3 – Toshiba

###### 3.2.5.3.1 – Satellite

3.2.6 – Desktops

3.2.6.1 – eMachines

3.2.6.2 – Igel

3.2.6.3 – Soekris

3.2.6.4 – IBM/Lenovo

3.2.6.4.1 – ThinkCentre

3.3 – Virtual Machines

3.3.1 – Qemu

3.3.1.1 – Installation

3.3.1.2 – Post-Installation Booting

3.3.1.2.1 – Multiboot

3.3.1.4 – Networking

3.3.1.4.1 – Linux VDE

3.3.1.4.2 – OpenBSD tun(4)

3.3.1.4.3 – Windows TAP

3.3.1.5 – Audio

3.3.1.6 – Graphics

3.3.2 – Virtualbox

3.3.2.1 – Ethernet

3.3.2.2 – Audio

3.3.2.3 – Graphics

3.3.2.4 – Known Working Versions

3.3.3 – Virtio

## **4 – 9front Installation Guide**

4.1 – Pre-installation Checklist

4.2 – Creating bootable 9front install media

4.2.1 – ISO image

4.2.2 – USB drive

4.2.2.1 – Creating on Plan 9

4.2.2.2 – Creating on Linux

4.2.2.3 – Bootargs

4.3 – Performing a simple install

4.3.1 – boot

4.3.2 – bootargs

4.3.3 – user

4.3.4 – vgasize, monitor, mouseport

4.3.4.1 – Changing screen resolution

4.3.5 – inst/start

4.3.6 – configfs

4.3.7 – partdisk

4.3.8 – prepdisk

4.3.9 – mountfs

4.3.10 – configdist

4.3.11 – confignet

4.3.11.1 – dhcp

4.3.11.2 – manual

4.3.12 – mountdist

4.3.13 – copydist

4.3.14 – sysname

4.3.15 – tzsetup

4.3.16 – bootsetup

4.3.17 – finish

4.4 – Encrypted Partitions



## **5 – Building the System from Source**

5.1 – Why should I build my system from source?

5.2 – Building 9front from source

5.2.1 – Update sources

5.2.1.1 – hgrc

5.2.2 – Building from source

5.2.2.1 – Cross compiling

5.3 – Building an ISO

5.4 – Common Problems when Compiling and Building

## **6 – Networking**

6.1 – Before we go any further

6.2 – Network configuration

6.2.1 – Host name

6.2.2 – Identifying and setting up your network interfaces

6.2.2.1 – Wifi

6.2.2.1.1 – Interfaces

6.2.2.1.1.1 – wavelan

6.2.2.1.1.2 – wavelanpci

6.2.2.1.1.3 – iwl

6.2.2.1.1.4 – rt2860

6.2.2.1.1.5 – wpi

6.2.2.1.2 – WPA

6.2.2.1.3 – Wifi Roaming

6.2.2.1.4 – Wifi Debug

6.2.3 – IP address

6.2.4 – Default gateway

6.2.5 – DNS Resolution

- 6.2.5.1 – Caching DNS server
- 6.2.5.2 – DNS authoritative name server
- 6.2.6 – Network-wide configuration
- 6.2.7 – Activating the changes
  - 6.2.7.1 – NIC
  - 6.2.7.2 – cs
  - 6.2.7.3 – dns
- 6.2.8 – Verifying network settings
  - 6.2.8.1 – Checking routes
    - 6.2.8.1.1 – Adding static routes
- 6.2.9 – Setting up your 9front box as a forwarding gateway
- 6.2.10 – Setting up aliases on an interface
- 6.3 – How do I filter and firewall with 9front?
- 6.4 – Dynamic Host Configuration Protocol (DHCP)
  - 6.4.1 – DHCP client
  - 6.4.2 – DHCP server
- 6.5 – PPP
- 6.6 – Setting up a network bridge in 9front
- 6.7 – How do I boot from the network?
  - 6.7.1 – How do I tcp boot?
    - 6.7.1.1 – Passing arguments to ipconfig at the bootargs prompt
  - 6.7.2 – How do I boot using PXE?

## **7 – System Management**

- 7.1 – Plan 9 Services Overview
  - 7.1.1 – What is the kernel?
  - 7.1.2 – What is the file server?
  - 7.1.3 – What is the auth server?

- 7.1.4 – What is the cpu server?
- 7.1.5 – What is a terminal?
- 7.2 – Kernel configuration and maintenance
  - 7.2.1 – How do I mount the 9fat partition?
  - 7.2.2 – How do I modify plan9.ini?
  - 7.2.3 – Kernel configuration file
  - 7.2.4 – Kernel drivers
  - 7.2.5 – How do I install a new kernel?
- 7.3 – Fileserver configuration and maintenance
  - 7.3.1 – Adding users
  - 7.3.2 – Configuring nvram
  - 7.3.3 – Setting up a listener for network connections
- 7.4 – Auth server configuration and maintenance
  - 7.4.1 – Configuring an auth server
    - 7.4.1.1 – Avoiding an ndb entry for the auth server
  - 7.4.2 – Adding users
  - 7.4.3 – secstored
  - 7.4.3 – Adding users to secstore
- 7.5 – Cpu server configuration and maintenance
  - 7.5.1 – Configuring a cpu server
- 7.6 – Terminal configuration and maintenance
  - 7.6.1 – Configuring a terminal
  - 7.6.2 – Configuring a Terminal to Accept cpu Connections
  - 7.6.3 – UTC Timesync
- 7.7 – Mail server configuration and maintenance
  - 7.7.1 – smtpd.conf
  - 7.7.2 – rewrite

7.7.3 – names.local

7.7.4 – remotemail

7.7.5 – SMTP over TLS

7.7.6 – IMAP4 over TLS

7.7.7 – Spam Filtering

7.7.7.1 – ratfs

7.7.7.2 – scanmail

7.8 – Web server configuration and maintenance

7.8.1 – ip/httpd

7.8.2 – rc-httpd

7.9 – TLS certificates

## **8 – Using 9front**

8.1 – rc

8.1.1 – Prompts

8.1.2 – /env

8.2 – rio

8.2.1 – The Pop-up Menu

8.2.2 – Window control

8.2.3 – Text in rio windows

8.2.4 – Scrolling

8.2.5 – Mouse Chording

8.2.6 – Keyboard Shortcuts

8.2.7 – Color scheme

8.2.8 – Why is rio like this?

8.2.9 – tips

8.2.9.1 – Taking a screenshot

8.3 – Text Editors

8.3.1 – sam

8.3.1.1 – Scrolling

8.3.1.2 – Mouse Chording

8.3.1.3 – Why does sam have a separate snarf buffer from rio?

8.3.1.4 – Keyboard Shortcuts

8.3.2 – acme

8.4 – Internet

8.4.1 – Mail

8.4.1.1 – upasfs

8.4.1.1.1 – Reading gmail via IMAP

8.4.1.2 – nedmail

8.4.1.3 – nupas

8.4.2 – NNTP

8.4.3 – IRC

8.4.3.1 – ircrc

8.4.3.2 – irc7

8.4.4 – FTP

8.4.5 – HTTP

8.4.5.1 – mothra

8.4.5.2 – abaco

8.4.5.3 – hget

8.4.5.4 – charon

8.4.5.5 – i

8.4.6 – SSH

8.4.6.1 – ssh

8.4.6.2 – ssh2

8.4.6.3 – scpu

- 8.4.6.4 – OpenSSH
- 8.4.6.5 – sftpfs
  - 8.4.6.5.1 – Mounting a remote u9fs share over SSH
- 8.4.7 – secstore
- 8.5 – Audio
- 8.6 – External Media
  - 8.6.1 – Mount an ISO9660 CD-ROM
  - 8.6.2 – Burn a CD-ROM
  - 8.6.3 – Mount a FAT formatted USB device
- 8.7 – Emulation
  - 8.7.1 – Linux Emulation
  - 8.7.2 – Nintendo
  - 8.7.3 – Sega
- 8.8 – Additional Software
  - 8.8.1 – 9front sources server
  - 8.8.2 – 9front contrib
- 8.9 – Bootstrapping architectures not included on the ISO
  - 8.9.1 – amd64
  - 8.9.2 – Raspberry Pi
- 8.10 – ACPI
  - 8.10.1 – Enabling ACPI
- 8.11 – MP Table Override
- 8.12 – Revision Control
  - 8.12.1 – cvs
  - 8.12.2 – git
  - 8.12.3 – Mercurial
  - 8.12.4 – svn




L - Languages





## FQA 0 – Introduction to Plan 9

### 0.1 – What is Plan 9?



Plan 9 is a research operating system from the same group who created UNIX at Bell Labs Computing Sciences Research Center (CSRC). It emerged in the late 1980s, and can be seen as an attempt to evolve some of the same ideas that informed UNIX into the era of networking and graphics. Rob Pike has described Plan 9 as "an argument" for simplicity and clarity, while others have described it as "UNIX, only moreso."

From *The Use of Name Spaces in Plan 9*:

Plan 9 argues that given a few carefully implemented abstractions it is possible to produce a small operating system that provides support for the largest

systems on a variety of architectures and networks.

From the `intro(1)` man page:

Plan 9 is a distributed computing environment assembled from separate machines acting as terminals, CPU servers, and file servers. A user works at a terminal, running a window system on a raster display. Some windows are connected to CPU servers; the intent is that heavy computing should be done in those windows but it is also possible to compute on the terminal. A separate file server provides file storage for terminals and CPU servers alike.

The two most important ideas in Plan 9 are:

- private namespaces (each process constructs a unique view of the hierarchical file system)
- file interfaces (familiar from UNIX, but taken to the extreme: all resources in Plan 9 look like file systems)

Most everything else in the system falls out of these two basic ideas.

Read: `intro(1)`; *Plan 9 from Bell Labs*; *Designing Plan 9*, originally delivered at the UKUUG Conference in London, July 1990; and *FQA 7 – System Management*; for a more detailed overview of Plan 9's design.



Today, Plan 9 continues in its original form, as well as in several derivatives and forks:

Plan 9 from Bell Labs — The original Plan 9. Effectively dead, all the developers have been run out of the Labs and/or are on display at Google.

Plan 9 from User Space — Plan 9 userspace ported/imitated for UNIX (specifically OS X)

9legacy — David du Colombier's cherry picked collection of patches from various

people/forks to Bell Labs Plan 9. (it is not a fork)

9atom — Erik Quanstrom's fork of Plan 9, maintained to erik's needs and occasionally pilfered by 9front

9front — (that's us) (we rule (we're the tunnel snakes))

NIX — High performance cloud computing is NIX — imploded in a cloud of political acrimony and retarded bureaucratic infighting

NxM — A kernel for manycore systems — never spotted in the wild

Clive — A new operating system from Francisco J. Ballesteros, designed to generate grantwriting practice material and research projects for otherwise indolent students

Akaros — Akaros is an open source, GPL-licensed operating system for manycore architectures. Has no bearing on anything but has attracted grant money.

Harvey — Harvey is an effort to get the Plan 9 code working with gcc and clang.

Inferno — Inferno is a distributed operating system started at Bell Labs, but is now developed and maintained by Vita Nuova Holdings as free software. Just kidding it is not developed or maintained.

### 0.1.1 – Plan 9 is not UNIX



In the words of the Bell Labs Plan 9 wiki:

Plan 9 is not Unix. If you think of it as Unix, you may become frustrated when something doesn't exist or when it works differently than you expected. If you think of it as Plan 9, however, you'll find that most of it works very smoothly, and that there are some truly neat ideas that make things much cleaner than you have seen before.

Confusion is compounded by the fact that many UNIX commands exist on Plan 9 and behave in similar ways. In fact, some of Plan 9's userland (such as the `upass` mail interface, the `sam` text editor, and the `rc` shell) are carried over directly from Research UNIX 10th Edition. Further investigation reveals that many ideas found in Plan 9 were explored in more primitive form in the later editions of Research UNIX.

However, Plan 9 is a completely new operating system that makes no attempt to conform to past prejudices. The point of the exercise (circa the late 1980s) was to avoid past problems and explore new territory.

Read: *UNIX to Plan 9 command translation*

### 0.1.1.1 – Plan 9 is not plan9port

Plan 9 from User Space (also known as plan9port or p9p) is a port of many Plan 9 from Bell Labs libraries and applications to UNIX-like operating systems. Currently it has been tested on a variety of operating systems including: Linux, Mac OS X, FreeBSD, NetBSD, OpenBSD, Solaris and SunOS.

Plan9port consists of a combination of mostly unaltered Plan 9 userland utilities and various attempts to imitate Plan 9 kernel interfaces using available UNIX programs and commands. Some of the imitations are more successful than others. In all, plan9port does not accurately represent the experience of using actual Plan 9.

### 0.1.1.2 – Plan 9 is not Inferno

	Inferno	Embedded Linux
Do you know the Designers	Yes, created by the inventors of the UNIX operating system and the C programming language	No
Designed for Embedded Systems	Yes	Server operating system modified for embedded systems
Get all the Source Code	Yes	Maybe
Supported Architectures	MIPS, s3b, StrongARM and PowerPC	x86 and i
Emulated Environments	Windows 95/98, Windows 2000/NT, Linux, etc. FreeBSD and others	None
Write Once, Run Anywhere	Yes	No
Unicode Character Set Support	No	Maybe
Rogity Free	Yes	Maybe
Protect your Intellectual Property	Yes	Maybe
Available Today	Yes	Maybe

The Network Operating System for distributed application development on network devices

VITA NUOVA  
Tel: 1-866-DANTEOS  
www.vitanuova.com

Inferno is a distributed operating system also started at Bell Labs, but is now developed and maintained by Vita Nuova Holdings as free software. It employs many ideas from Plan 9 (and even shares some source code), but is a completely different OS.

**Note:** Inferno shares some compatible interfaces with Plan 9, including the 9P/Styx protocol.

## 0.1.2 – Plan 9 is not a product

Path: utzoo!utgpu!water!watmath!clyde!bellcore!faline!thumper!ulysses!smb  
From: s...@ulysses.homer.nj.att.com (Steven Bellovin)  
Newsgroups: comp.unix.wizards  
Subject: Re: Plan 9? (+ others)  
Message-ID: <10533@ulysses.homer.nj.att.com>  
Date: 23 Aug 88 16:19:40 GMT  
References: <846@yunexus.UUCP> <282@umbio.MIAMI.EDU> <848@yunexus.UUCP>  
Organization: AT&T Bell Laboratories, Murray Hill  
Lines: 33

“Plan 9” is not a product, and is not intended to be. It is research -- an experimental investigation into a different way of computing. The developers started from several basic assumptions: that CPUs are very cheap but that we don't really know how to combine them effectively; that \*good\* networking is very important; that an intelligent user interface (complete with dot-mapped display and mouse) is a Right Decision; that existing systems with networks, mice, etc., are not the correct way to do things, and in particular that today's workstations are not the way to go. (No, I won't bother to explain all their reasoning; that's a long and separate article.) Finally, the UNIX system per se is dead as a vehicle for serious research into operating system structure; it has grown too large, and is too constrained by 15+ years of history.

Now -- given those assumptions, they decided to throw away what we have today and design a new system. Compatibility isn't an issue -- they are not in the product-building business. (Nor are they in the “let's make another clever hack” business.) Of course aspects of Plan 9 resemble the UNIX system quite strongly -- is it any surprise that Pike, Thompson, et al., think that that's a decent model to follow? But Plan 9 isn't, and is not meant to be, a re-implementation of the UNIX system. If you want, call it a UNIX-like system.

Will Plan 9 ever be released? I have no idea. Will it remain buried? I hope not. Large companies do not sponsor large research organizations just for the prestige; they hope for an (eventual) concrete return in the form of concepts that can be made into (or incorporated into) products.

--Steve Bellovin

Disclaimer: this article is not, of course, an official statement from AT&T. Nor is it an official statement of the reasoning behind Plan 9. I do think it's accurate, though, and I'm sure I'll be told if I'm wrong...

## 0.1.3 – Plan 9 is not for you

**Let's be perfectly honest.** Many features that today's "computer experts" consider to be essential to computing (javascript, CSS, HTML5, etc.) either did not exist when Plan 9 was abandoned, or were purposely left out of the operating system. You might find this to be an unacceptable obstacle to adopting Plan 9 into your daily workflow. If cannot

imagine a use for a computer that does not involve a web browser, Plan 9 may not be for you.

See: <http://harmful.cat-v.org/software/>

## 0.2 – Why Plan 9?

You may ask yourself, well, how did I get here? In the words of Plan 9 contributor Russ Cox:

Why Plan 9 indeed. Isn't Plan 9 just another Unix clone? Who cares?

Plan 9 presents a consistent and easy to use interface. Once you've settled in, there are very few surprises here. After I switched to Linux from Windows 3.1, I noticed all manner of inconsistent behavior in Windows 3.1 that Linux did not have. Switching to Plan 9 from Linux highlighted just as much in Linux.

One reason Plan 9 can do this is that the Plan 9 group has had the luxury of having an entire system, so problems can be fixed and features added where they belong, rather than where they can be. For example, there is no tty driver in the kernel. The window system handles the nuances of terminal input.

If Plan 9 was just a really clean Unix clone, it might be worth using, or it might not. The neat things start happening with user-level file servers and per-process namespaces. In Unix, `/dev/tty` refers to the current window's output device, and means different things to different processes. This is a special hack enabled by the kernel for a single file. Plan 9 provides full-blown per-process namespaces. In Plan 9 `/dev/cons` also refers to the current window's output device, and means different things to different processes, but the window system (or telnet daemon, or ssh daemon, or whatever) arranges this, and does the same for `/dev/mouse`, `/dev/text` (the contents of the current window), etc.

Since pieces of file tree can be provided by user-level servers, the kernel need not know about things like DOS's FAT file system or GNU/Linux's EXT2 file system or NFS, etc. Instead, user-level servers provide this functionality when desired. In Plan 9, even FTP is provided as a file server: you run `ftpf`s and the files on the server appear in `/n/ftp`.

We need not stop at physical file systems, though. Other file servers synthesize files that represent other resources. For example, `upas/fs` presents your mail box as a file tree at `/mail/fs/mbox`. This models the recursive structure of MIME messages especially well.

As another example, `cdfs` presents an audio or data CD as a file system, one file per track. If it's a writable CD, copying new files into the `/mnt/cd/wa` or `/mnt/cd/wd` directories does create new audio or data tracks. Want to fixate the CD as audio or data? Remove one of the directories.

Plan 9 fits well with a networked environment, files and directory trees can be imported from other machines, and all resources are files or directory trees,

it's easy to share resources. Want to use a different machine's sound card? Import its /dev/audio. Want to debug processes that run on another machine? Import its /proc. Want to use a network interface on another machine? Import its /net. And so on.

Russ Cox

### **0.2.1 – What do people like about Plan 9?**

Descriptive testimony by long time Plan 9 users Charles Forstyh, Anthony Sorace and Geoff Collyer:

[http://www.plan9.bell-labs.com/wiki/plan9/what\\_do\\_people\\_like\\_about\\_plan\\_9/](http://www.plan9.bell-labs.com/wiki/plan9/what_do_people_like_about_plan_9/)

#### **0.2.1.1 – What do you use Plan 9 for?**

if you dont know what this is usefull

then its not for you

See: *FQA 8 – Using 9front*

### **0.2.2 – What do people hate about Plan 9?**

John floren provides a humorous(?) overview of a typical new user's reactions to Plan 9:

Hi! I'm new to Plan 9. I'm really excited to work with this new Linux system.

I hit some questions.

- 1 How do I run X11?
- 2 Where is Emacs?
- 3 The code is weird. It doesn't look like GNU C at all. Did the people who wrote Plan 9 know about C?
- 4 I tried to run mozilla but it did not work. How come?

Is this guy you?

See: BUD/S

#### **0.2.2.1 – What is not in Plan 9**

A summary of common features you may have been expecting that are missing from Plan 9:

<http://c2.com/cgi/wiki?WhatIsNotInPlanNine>

### 0.2.3 – Why did Plan 9's creators give up on Plan 9?



All of the people who worked on Plan 9 have moved on from Bell Labs and/or no longer work on Plan 9. Various reasons have been articulated by various people.

Russ Cox:

I ran Plan 9 from Bell Labs as my day to day work environment until around 2002. By then two facts were painfully clear. First, the Internet was here to stay; and second, Plan 9 had no hope of keeping up with web browsers. Porting Mozilla to Plan 9 was far too much work, so instead I ported almost all the Plan 9 user level software to FreeBSD, Linux, and OS X.

Russ Cox (again):

The standard set up for a Plan 9 aficionado here seems to be a Mac or Linux machine running Plan 9 from User Space to get at sam, acme, and the other tools. Rob, Ken, Dave, and I use Macs as our desktop machines, but we're a bit of an exception. Most Google engineers use Linux machines, and I know of quite a few ex-Bell Labs people who are happy to be using sam or acme on those machines. My own setup is two screens. The first is a standard Mac desktop with non-Plan 9 apps and a handful of 9terms, and the second is a full-screen acme for getting work done. On Linux I do the same but the first screen is a Linux desktop running rio (formerly dhog's 8½).

More broadly, every few months I tend to get an email from someone who is happy to have just discovered that sam is still maintained and available for modern systems. A lot of the time these are people who only used sam on Unix, never on Plan 9. The plan9port.tgz file was downloaded from 2,522 unique IP addresses in 2009, which I suspect is many more than Plan 9 itself. In that sense, it's really nice to see the tools getting a much wider exposure than they used to.

I haven't logged into a real Plan 9 system in many years, but I use 9vx



occasionally when I want to remind myself how a real Plan 9 tool worked. It's always nice to be back, however briefly.

Russ

Russ Cox continues:

- > Can you briefly tell us why you (Russ, Rob, Ken and Dave)
- > no longer use Plan9 ?
- > Because of missing apps or because of missing driver for your hardware ?
- > And do you still use venti ?

Operating systems and programming languages have strong network effects: it helps to use the same system that everyone around you is using. In my group at MIT, that meant FreeBSD and C++. I ran Plan 9 for the first few years I was at MIT but gave up, because the lack of a shared system made it too hard to collaborate. When I switched to FreeBSD, I ported all the Plan 9 libraries and tools so I could keep the rest of the user experience.

I still use venti, in that I still maintain the venti server that takes care of backups for my old group at MIT. It uses the plan9port venti, vbackup, and vnfs, all running on FreeBSD. The venti server itself was my last real Plan 9 installation. It's Coraid hardware, but I stripped the software and had installed my own Plan 9 kernel to run venti on it directly. But before I left MIT, the last thing I did was reinstall the machine using FreeBSD so that others could help keep it up to date.

If I wasn't interacting with anyone else it'd be nice to keep using Plan 9. But it's also nice to be able to use off the shelf software instead of reinventing wheels (9fans runs on Linux) and to have good hardware support done by other people (I can shut my laptop and it goes to sleep, and even better, when I open it again, it wakes up!). Being able to get those things and still keep most of the Plan 9 user experience by running Plan 9 from User Space is a compromise, but one that works well for me.

Russ

Rob Pike:

What Russ says is true but for me it was simpler. I used Plan 9 as my local operating system for a year or so after joining Google, but it was just too inconvenient to live on a machine without a C++ compiler, without good NFS and SSH support, and especially without a web browser. I switched to Linux but found it very buggy (the main problem was most likely a bad graphics board and/or driver, but still) and my main collaborator (Robert Griesemer) had done the ground work to get a Mac working as a primary machine inside Google, and Russ had plan9port up, so I pushed plan9port onto the Mac and have been there ever since, quite happily. Nowadays Apples are officially supported so it's become easy, workwise.

I miss a lot of what Plan 9 did for me, but the concerns at work override that.

-rob

## 0.2.4 – What is the deal with Plan 9's weird license?



Over the years Plan 9 has been released under various licenses, to the consternation of many.

The first edition, released in 1992, was made available only to universities. The process for acquiring the software was convoluted and prone to clerical error. Many potential users had trouble obtaining it within a reasonable time frame and many complaints were voiced on the eventual Plan 9 Internet mailing list.

The second edition, released in 1995 in book-and-CD form under a relatively standard commercial license, was available via mailorder as well as through a special telephone number for a price of approximately \$350 USD.

### 0.2.4.1 – Richard Stallman hates the Plan Nine license (circa 2000)

In the year 2000, the third edition of Plan 9 was finally released under a custom "open source" license, the Plan 9 License. Richard Stallman was not impressed:

When I saw the announcement that the Plan Nine software had been released as "open source", I wondered whether it might be free software as well. After studying the license, my conclusion was that it is not free; the license contains several restrictions that are totally unacceptable for the Free Software Movement. (See <http://www.gnu.org/philosophy/free-sw.html>.)

Read more here:

<http://www.linuxtoday.com/developer/2000070200704OPLFSW>

#### 0.2.4.2 – Theo de Raadt hates the Plan 9 license (circa 2003)

In the year 2002, the fourth edition of Plan 9 was released under the Lucent Public License. Theo de Raadt was not impressed:

The new license is utterly unacceptable for use in a BSD project.

Actually, I am astounded that the OSI would declare such a license acceptable.

That is not a license which makes it free. It is a \*contract\* with consequences; let me be clear -- it is a contract with consequences that I am unwilling to accept.

Read more here:

<http://9fans.net/archive/2003/06/270>

#### 0.2.4.3 – Everyone hates the Plan 9 license (circa 2014)

In 2014, portions of the Plan 9 source code were again re-licensed, this time under the GPLv2, for distribution with the University of California, Berkeley's Akaros operating system. Various parties were not impressed.

Russ Cox tried to make sense of the situation by commenting in a Hacker News thread:

When you ask "why did big company X make strange choice Y regarding licensing or IP", 99 times out of 100 the answer is "lawyers". If the Plan 9 group had had its way, Plan 9 would have been released for free under a trivial MIT-like license (the one used for other pieces of code, like the one true awk) in 2003 instead of creating the Lucent Public License. Or in 2000 instead of creating the "Plan 9 License". Or in 1995 instead of as a \$350 book+CD that came with a license for use by an entire "organization". Or in 1992 instead of being a limited academic release.

Thankfully I am not at Lucent anymore and am not privy to the tortured negotiations that ended up at the obviously inelegant compromise of "The University of California, Berkeley, has been authorised by Alcatel-Lucent to release all Plan 9 software previously governed by the Lucent Public License, Version 1.02 under the GNU General Public License, Version 2." But the odds are overwhelming that the one-word answer is "lawyers".

Some have suggested that confusion about licensing has contributed to Plan 9's failure to supplant UNIX in the wider computing world.

### 0.3 – Further Reading

#### 0.3.1 – Plan 9 papers

Academic papers that describe the Plan 9 operating system are available here:

[http://doc.cat-v.org/plan\\_9/](http://doc.cat-v.org/plan_9/)

### 0.3.2 – Man pages

Section (1) for general publicly accessible commands.

Section (2) for library functions, including system calls.

Section (3) for kernel devices (accessed via bind(1)).

Section (4) for file services (accessed via mount).

Section (5) for the Plan 9 file protocol.

Section (6) for file formats.

Section (7) for databases and database access programs.

Section (8) for things related to administering Plan 9.

### 0.3.3 – Web pages

The official website for the Plan 9 project is located at: <http://plan9.bell-labs.com/plan9>

Bell Labs Plan 9 wiki: [http://www.plan9.bell-labs.com/wiki/plan9/plan\\_9\\_wiki](http://www.plan9.bell-labs.com/wiki/plan9/plan_9_wiki)

A lot of valuable information can be found at <http://cat-v.org> regarding aspects of UNIX, Plan 9 and software in general.

### 0.3.4 – Books



*Introduction to OS Abstractions Using Plan 9 From Bell Labs*, by Francisco J Ballestros (nemo)

*Notes on the Plan 9 3rd Edition Kernel*, by Francisco J Ballestros (nemo)

*The UNIX Programming Environment*, by Brian W. Kernighan (bwk) and Rob Pike (rob)  
(this book is the most clear, concise and eloquent expression of the Unix and 'tool'  
philosophies to date)

*9front Frequently Questioned Answers*, by Stanley Lieber (sl) (the document you are  
reading right now, in book form)

## FQA 1 – Introduction to 9front



### 1.1 – What is 9front?

Plan9front (or 9front) is a fork of the Plan 9 from Bell Labs operating system. The project was started to remedy a perceived lack of devoted development resources inside Bell Labs,[citation needed] and has accumulated various fixes and improvements.

This FQA specifically covers only the most recent release of 9front.

#### 1.1.1 – Cirno's been our mascot for a while.

**Our mascot**

Glenda is getting old, Cirno is now our mascot.




<http://touhou.wikia.com/wiki/Cirno>

### Pro

- girl
- has magical powers
- associated with 9
- upsets kfx
- she is known to be the strongest

## Alternatives



### 1.2 – On what systems does 9front run?

9front runs on the following platforms:

- 386
- amd64
- arm

Read: *FQA 3.2 – Selecting Hardware*



### 1.3 – Why might I want to use 9front?

New users frequently want to know whether 9front is superior to some other free UNIX-like operating system. That question is largely unanswerable and is the subject of countless (and useless) religious debates. Do not, under any circumstances, ask such a question on IRC or on a 9front mailing list.

Whether 9front is right for you is a question that only you can answer.

#### 1.3.1 – New Features

- `/shr`, global mountpoint device `shr(3)`
- `/mnt` is provided by `mntgen(4)`
- `#A`, audio drivers for sb16, intel hd audio and ac97 (both playback and recording supported!) `audio(3)`
- New BIOS based boot loader `9boot(8)` featuring a console and support for FAT/ISO/PXE and being small (<8K)
- New EFI based boot loader `efi`
- Made kernel compliant to multiboot specification so it can be booted by `qemu` or `grub` directly.
- Interruptable kernel qlocks (`eqlock`)
- Deferred clunks (`closeproc`) for cached mounts
- New `rc` based `boot(8)` allows breaking into a shell at any time
- Default file system is an improved `cwfs(4)` (`cwfs64x`)
- Support for encrypted partitions `cryptsetup(8)`. Read: FQA 4 – 9front Installation Guide
- New screen fonts: `dejavu`, `germgoth`, `vga`
- No central `replica`; source updates are done with `hg(1)` (Mercurial)
- Keyboard events with `/dev/kbd`. Read: `kbd(8)` and `rio(4)`
- `/lib/ken`, `/lib/rob`, `/lib/rsc`
- New `listen(8)` `-p maxprocs` option
- Always available network `aan(8)` support in `cpu(1)`
- MSI (message signalled interrupts), avoids problems with broken MP tables. Read: `icanhasmsi(8)` (dump MSI info)
- Legacy free ACPI support (aml interpreter `libaml`, mp interrupt routing,

scram)

- New `rio(1)` `-b` option (black window backgrounds)
- USB CD-ROM boot/install
- USB drive boot
- Improved USB mouse support
- Support for USB ptp cameras
- Stable across machines USB device names
- VGA initialization done by interpreting the VESA BIOS with `realemu(8)`, working VESA screen blanking.
- `/dev/kbd` and clipboard charset support for `vnc(1)`
- New `webfs(4)` with HTTP1.1 and Keep-Alive support.
- Qemu/KVM virtio block device and ethernet drivers. Read: *FQA 4.5.1.3 – Virtio*
- Mouse wheel and chording support in `sam(1)`
- Elliptic curve cryptography `ec(2)`
- Working interrupt key (Del) in console
- Wifi support with wpa/wpa2
- SSE support
- System-wide support for internationalized domain names
- Unicode support in `vt(1)`
- `pc64`, kernel for amd64

### 1.3.1.1 – New Programs

- `"` and `" "` (print, repeat previous command)
- `alarm(1)` – timeouts in `rc` scripts
- `audio(1)` – mp3, ogg, flac,  $\mu$ law, wav
- `bullshit(1)` – print out a stream of bullshit
- `cifsd(8)` – CIFS/SMB server
- `derp(1)` – find changes between directories

- `feminize(1)` — replace sexist remarks
- `fplot(1)`
- New games: `doom`, `glendy`, `mandel`, `mines`, `mole`, `packet`
- `hg(1)` and `hgfs(4)` (Mercurial)
- `hget(1)` — rewritten in `rc`, now uses `webfs`
- `hpost(1)` — extract and post HTML forms
- `hold(1)` — simple text editor
- `ipserv(8)` — proxy servers `socksd` and `hproxy`
- `memory(1)` — check memory usage
- `mothra(1)` — Tom Duff's web browser, now uses `webfs`
- `netaudit(8)` — network configuration checker
- `newt(1)` — Usenet client
- `nietzsche(1)` — print out Nietzsche quote
- `nintendo(1)` — Nintendo emulators: `gb`, `gba`, `nes`, `snes`
- `page(1)` — zoom and enhance!
- `paint(1)` — drawing program
- `play(1)` — audio player
- `pstree(1)` — print tree-like map of current processes and sub-processes
- `rotate(1)` — rotate or mirror a picture
- `resize(1)` — fast but low quality image resampler
- `scram(8)` — ACPI and APM shutdown
- `sega(1)` — Sega Megadrive/Genesis emulator: `md`
- `ssam(1)` — stream interface to `sam`
- `sysinfo(1)` — print hardware report
- `sysupdate(1)` — update the local `hg` repository
- `theo(1)` — print out insults from Theo de Raadt
- `tput(1)` — measure read throughput

- `troll(1)` — automated trolling
- `tap(1)` — follow the pipes of a process
- `tif(1)` — tiff decoder
- `tojpg(1)` — jpeg encoder
- `totif(1)` — tiff encoder
- `torrent(1)` — bittorrent client

### 1.3.1.2 – New Hardware Support

#### Audio

- AC97 — `/sys/src/9/pc/audioac97.c`
- HDA — `/sys/src/9/pc/audiohda.c`
- SB 16/ESS — `/sys/src/9/pc/audiosb16.c`

#### Ethernet

- ADMtek Pegasus — `/sys/src/cmd/nusb/ether/ae.c`
- Broadcom BCM57xx — `/sys/src/9/pc/etherbcm.c`
- Realtek RTL8150 — `/sys/src/cmd/nusb/ether/ur1.c`

#### Wifi

- Intel Centrino Advanced-N 6205 (iwl-6005)
- Intel Centrino Ultimate-N (iwl-6000)
- Intel Centrino Wireless-N 100
- Intel WiFi Link 1000/4965/5100/5300 AGN
- Intel PRO Wireless 3945ABG (wpi-3945abg)
- Ralink RT2860/RT3090

## Tablets

- Wacom WACF004, WACF008

`/sys/src/cmd/aux/wacom.c`, `/sys/src/cmd/aux/tablet.c`

## Video

- AMD Geode LX driver

`/sys/src/cmd/aux/vga/geode.c` `/sys/src/9/pc/vgageode.c`

- Intel G45 and Ivy Bridge driver

`/sys/src/cmd/aux/vga/igfx.c` `/sys/src/9/pc/vgaigfx.c`

## SD Card

- Ricoh — `/sys/src/9/pc/pmmc.c`

Read: *FQA 3.2 – Known Working Hardware* for a list of complete machines known to work.



## 1.4 – Is 9front really free?

Yes.

Read: *FQA 0.2.4 – What is the deal with Plan 9's weird license?*

**Note:** The maintainers of 9front support the project largely from their own pockets. This includes the time spent programming for the project, equipment used, the network resources used to distribute 9front to you, and the time spent answering questions and investigating users' bug reports. The 9front developers are not independently wealthy and even small contributions of time, equipment, and resources make a big difference.

### **1.5 – How can I help support 9front?**

We are greatly indebted to the people and organizations that have contributed to the 9front project.

### **1.6 – Who maintains 9front?**

9front is maintained by a development team spread across many different countries.

### **1.7 – When is the next release of 9front?**

The 9front team makes new releases on a regular, but unscheduled, basis. More information on the development cycle can be found here:

<http://lists.9front.org>

### **1.8 – What is included with 9front?**

Some useful programs included with the operating system are:

- `acid` — Programmable symbolic debugger.
- `acme` — Text editor, window system, mail client and more.
- `ape` — ANSI/POSIX environment.
- `cwfs64x` — Cached-worm file server based on the original Ken's fs.
- `doom` — Science fiction horror-themed first-person shooter video game by id Software.
- `gs` — Aladdin Ghostscript (PostScript and PDF language interpreter).
- `hjfs` — A new, experimental fs.
- `mercurial` — Distributed source control management tool.
- `mk` — Tool for describing and maintaining dependencies between files.
- `mothra` — Web browser by Tom Duff.
- `newt` — NNTP client.
- `nintendo` — Nintendo Game Boy, NES, SNES and GBA emulators.

- `openssl` — Toolkit implementing the Secure Sockets Layer (SSL v2/v3) and Transport Layer Security (TLS v1) protocols as well as a general purpose cryptography library.
- `paint` — Drawing program.
- `page` — FAX, image, graphic, PostScript, PDF, epub, cbz viewer.
- `play` — Flac, ogg, mp3, sun, wav player.
- `plumber` — Mechanism for inter-process communication.
- `python 2.5.1` — Interpreted programming language.
- `rc` — Shell by Tom Duff.
- `rc-httpd` — Web server written in rc.
- `rio` — Rectangle multiplexer/window system.
- `sam` — Text editor.
- `sega` — Megadrive/Genesis emulator.
- `torrent` — BitTorrent client.
- `troff` — Text processor/typesetter.
- `upas` — A simpler approach to network mail.

## 1.9 – Can I use 9front as a desktop system?

This question is often asked in exactly this manner—with no explanation of what the asker means by "desktop". The only person who can answer that question is you, as it depends on what your needs and expectations are.

## 1.10 – Why is/isn't ProductX included?

People often ask why a particular product is or isn't included with 9front. The answer is based on two things: the wishes of the developers and compatibility with the goals of the project. A product will not be included simply because it is "nifty."

A major consideration is the wishes of the developers. The 9front developers are the ultimate judges of what does and doesn't go into the project. Just because an application is "good" doesn't mean the 9front project wishes to devote the resources needed to maintaining it, or that they will share other's enthusiasm about its place in 9front.

Incredibly, many users still see fit to demand otherwise.

## 1.11 – Fine, where can I get 9front?



If you simply cannot be dissuaded from trying 9front for yourself, obtain installation media from the following locations.

### 1.11.1 – Mirrors

#### 1.11.1.1 – 9front.iso

<http://9front.org/iso/>

<http://r-36.net/9front/>

<http://9front.r0x0r.me/>

<http://mirror.9fags.net/>

[http://9hal.ath.cx/usr/cinap\\_lenrek/9front.torrent](http://9hal.ath.cx/usr/cinap_lenrek/9front.torrent)

[http://thefront.felloff.net/usr/cinap\\_lenrek/9front.torrent](http://thefront.felloff.net/usr/cinap_lenrek/9front.torrent)



Good Luck



## FQA 2 – Getting To Know 9front



### 2.1 – Web Pages

The official websites for the 9front project are located at:

- <http://9front.org> — main website
- <http://fqa.9front.org> — official documentation, frequently questioned answers
- <http://lists.9front.org> — mailing lists
- <http://man.9front.org> — man pages
- <http://wiki.9front.org> — informal documentation, developer scratchpad

Somewhat outdated, and not entirely applicable to 9front is the Bell Labs Plan 9 wiki:  
[http://www.plan9.bell-labs.com/wiki/plan9/plan\\_9\\_wiki](http://www.plan9.bell-labs.com/wiki/plan9/plan_9_wiki)

Some users have set up sites and pages with 9front specific information. As with everything on the Internet, a good search engine is going to make your life easier, as will a healthy dose of skepticism. As always, do not blindly enter commands you do not understand into your computer.

A lot of valuable information can be found at <http://cat-v.org> regarding aspects of UNIX, Plan 9 and software in general.

### 2.2 – Mailing Lists

The 9front project maintains several mailing lists which users should subscribe to and follow. To subscribe to a mailing list, send an e-mail message to [list-owner@9front.org](mailto:list-owner@9front.org),

where *list* is the name of the mailing list. That address is an automated subscription service. In the body of your message, on a single line, you should include the subscrib command. For example:

```
subscribe
```

After subscribing, list messages are received from and sent to *list@9front.org*, where *list* is the name of the mailing list.

The *lists* are as follows:

9front — general discussion

9front-bugs — report bugs

9front-commits — read-only log of commits to the 9front mercurial repository

9front-sysinfo — read-only log of `sysinfo(1)` output from user systems

To unsubscribe from a list, you will again send an e-mail message to *list-owner@9front.org*. It should look like this:

```
unsubscribe
```

Before posting a question on the 9front mailing list, please check the rest of this FQA. If asking a question possibly related to hardware, always include the output of `sysinfo(1)`!

## 2.2.1 – Mailing List Archives

Mailing list archives are available via 9p:

```
term% 9fs 9front.org
post...
term% ls /n/9front.org/lists
/n/9front.org/lists/9fans
/n/9front.org/lists/9fans.old
/n/9front.org/lists/9front
/n/9front.org/lists/9front-bugs
/n/9front.org/lists/9front-commits
/n/9front.org/lists/9front-sysinfo
/n/9front.org/lists/acme-sac
/n/9front.org/lists/cat-v
/n/9front.org/lists/inferno
/n/9front.org/lists/sam-fans
/n/9front.org/lists/werc
```

Each of these directories contains messages sent to the respective mailing list in `mdir` format, one message per file. Erik Quanstrom's `nupas upas/fs` can mount these directories as mailboxes, for reading directly with a mail client.

## 2.2.2 – Other useful mailing lists



9fans — Official mailing list of Plan 9 from Bell Labs.

9changes — Notifications of changes to Plan 9 from Bell Labs (unofficial).

9nag — Ping alerts for plan9.bell-labs.com (unofficial).

## 2.3 – Manual Pages

9front comes with extensive documentation in the form of manual pages. Considerable effort is made to make sure the man pages are up-to-date and accurate. In all cases, 1.) the source, followed by 2.) the man pages, are considered the authoritative source of information for 9front.

You can find all the 9front man pages on the web at the following sites:

<http://man.9front.org>

<http://9hal.ath.cx/sys/man>

<http://man.aiju.de>

<http://man.cat-v.org/9front>

<http://plan9.stanleylieber.com/man>

as well as on your 9front computer.

In general, if you know the name of a command or a manual page, you can read it by executing `man` command. For example: `man vi` to read about the MIPS binary emulator. If you don't know the name of the command, or if `man` command doesn't find the manual page, you can search the manual page database by executing `lookman something`, where `something` is a likely word that might appear in the title of the manual page you're looking for. For example:

```
term% lookman vesa
man 3 vga # vga(3)
man 8 realemu # realemu(8)
man 8 vga # vga(8)
```

The resulting list is presented in the form of commands that can be highlighted and sent using rio's mouse button 2 menu.

For many, having a hardcopy of the man page can be useful. To make a printable copy of a man page:

```
man -t vga realemu | lp -dstout > vga.ps
```

or:

```
man -t vga realemu | lp -dstout | ps2pdf > vga.pdf
```

### 2.3.1 – How do I display a man page source file?

The command `man -w` prints the location of the man page's source file:

```
term% man -w vga realemu
/sys/man/3/vga
/sys/man/8/realemu
/sys/man/8/vga
```

### 2.3.2 – How do I write my own manual page?

As was mentioned in *FQA 0.3.2 – Man pages*, Plan 9 man pages are divided into sections based upon the features they describe. For example, the troff source of the `io(1)` man page is located in section 1, in the file `/sys/man/1/io`. You can examine this troff source for a simple example of how to format and construct a typical man page.

For additional information, read: *How to Write a Plan 9 Manual Page* `man(6)`, `troff(1)`

## 2.4 – Reporting Bugs

Before crying "Bug!", please make sure that is really what you are dealing with. If instead, you are not understanding how something is done in 9front or how it works, and can't find out how to resolve the problem using the manual pages or *FQA 9 – Troubleshooting*, join `#9front` on `irc.freenode.net` or use the mailing list (`9front@9front.org`) to request help. If this is your first 9front experience, be realistic: you probably don't know what you're looking at and it is likely you did not discover an unknown bug. Also note that faulty hardware can mimic a software bug, please verify the current condition of your hardware before deciding you have found a "bug."

Remember, before submitting any bug report, please read *FQA 8 – Using 9front* and *FQA 9 – Troubleshooting*.



*"I have with Plan 9 Is there a way to"*

Proper bug reporting is one of the most important responsibilities of end users. Very detailed information is required to diagnose most serious bugs. Developers frequently get bugs reports via e-mail such as this:

From: joeuser@example.com  
To: 9front@9front.org  
Subject: HELP!!!

I have a PC and it won't boot!!!! It's a 486!!!!

Hopefully most people understand why such reports get summarily deleted. All bug reports should contain detailed information. If Joe User had really expected someone to help find this bug, he or she would have supplied more information... something like this:

From: smartuser@example.com  
To: 9front@9front.org  
Subject: Acme suicide: fault read.

I snarfed some text from rc and when I clicked 'paste' inside Acme, Acme crashed. See the attached file for the error message.

The output from sysinfo is also attached.

**Note:** Detailed information about your hardware is necessary if you think the bug *could be in any way* related to your hardware or hardware configuration.

## 2.4.1 – Your bug report SUCKS

A minimum useful bug report contains the following:

- A complete description of what you did
- A complete description of what you expected to happen
- A complete description of what happened, instead

Extra credit:

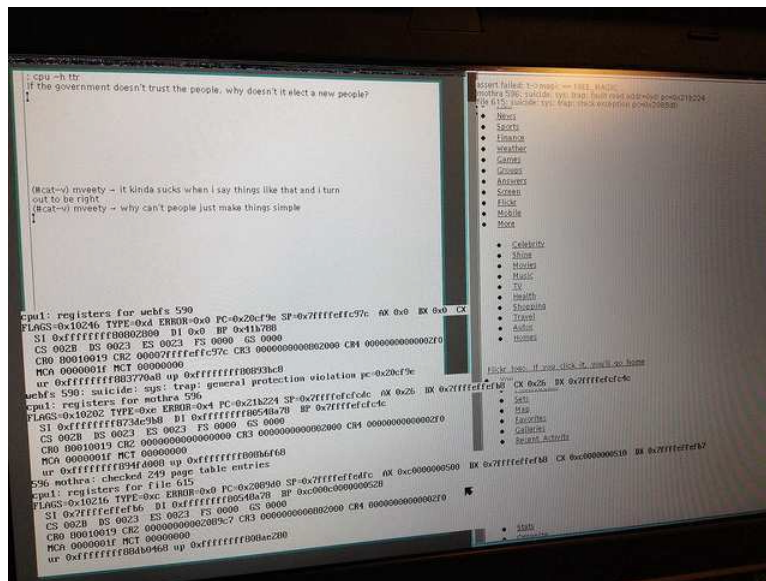
- The make and model of your computer
- The URL returned when you run `sysinfo -p` on the affected system

## 2.4.2 – How do I get more useful info for developers?

The following sub-sections offer some useful tips.

### 2.4.2.1 – Take a photo of the panic screen

Under some circumstances, you may not be able to recover text directly from a crashed system. In this situation, snap a photo of the screen, including as much of the screen in good focus as possible, and upload the image somewhere where it can be retrieved by developers.



### 2.4.2.2 – Make a stack trace of a crashed process

```
term% hget http://www.jwz.org/blog/feed | rssread
panic: D2B called on non-block c580 (double-free?)
rssread 882285: suicide: sys: trap: fault read addr=0x0 pc=0x000038ef
term% acid 882285
/proc/882285/text:386 plan 9 executable
/sys/lib/acid/port
/sys/lib/acid/386
acid: lstk()
abort()+0x0 /sys/src/libc/9sys/abort.c:6
ppanic(p=0xb5a0,fmt=0xbf72)+0x146 /sys/src/libc/port/malloc.c:166
    pv=0xc100
    msg=0xc310
    v=0xdffffee90
    n=0x2b
D2B(p=0xb5a0,v=0xc580)+0x5a /sys/src/libc/port/pool.c:966
    a=0xc578
poolfree1(v=0xc580,p=0xb5a0)+0x20 /sys/src/libc/port/pool.c:1190
    ab=0xc100
poolfree(p=0xb5a0,v=0xc580)+0x41 /sys/src/libc/port/pool.c:1325
free(v=0xc588)+0x23 /sys/src/libc/port/malloc.c:250
nextxmlpull(x=0xc548)+0x334 /usr/sl/src/xmlpull/xmlpull.c:221
nextxmlpull(x=0xc548)+0x2cb /usr/sl/src/xmlpull/xmlpull.c:316
main()+0x34 /usr/sl/src/rssread/rssread.c:159
    st=0x1
    f=0xc5a8
    r=0x0
    x=0xc548
_main+0x31 /sys/src/libc/386/main9.s:16
acid:
```



## FQA 3 – Hardware



### 3.1 – Selecting Hardware

Selecting appropriate hardware to run your 9front system on is important, as it can mean the difference between success and failure of a project.

If you are shopping for a new PC, whether you are buying it piece by piece or completely pre-built, you want to make sure first that you are buying reliable parts. In the PC world, this is not easy. **Bad or otherwise unreliable or mismatched parts can make 9front run poorly and crash often.** The best advice we can give is to be careful, and buy brands and parts that have been reviewed by an authority you trust. Sometimes, a higher-price machine is a better quality machine. Other times, it is simply more expensive.

There are certain things that will help bring out the maximum performance of your system:

**Let the application choose the hardware:** It is usually better to make an adjustment to the hardware you were planning on using rather than compromising on your application design because you have something you "really wanted to use."

**Identify your bottlenecks:** Don't pay extra for the cutting-edge processor if your application is restricted by disk I/O. Don't pay for fast disk if your system is restricted by network speed. Don't pay for much of anything if your bottleneck is a 300bps dial-up line.

**Keep it simple:** Simple hardware usually has simple problems. Complex hardware that isn't supposed to ever break may take you a long time to repair when it breaks anyway.

**Use hardware you understand, or learn the new hardware before you implement a production system:** Regardless of the technical merits of the

hardware, committing to use a particular type of hardware before you have become familiar with it, both how it works and how it fails, is foolish.

**Avoid cheap network adapters:** 9front supports a plethora of cheap network adapters. These adapters work great in home systems, and low or moderate throughput business and research environments. But, if you need high throughput and low impact on your server, you are better off buying a quality network adapter. Unfortunately, some name-brand adapters are not much better than the cheap adapters, and some potentially good adapters do not have accurate documentation available to write good drivers. Gigabit adapters often perform better than 10Mbps/100Mbps adapters, even when used on slower speed networks, due to superior buffering.

Check *FQA 3.2 – Known Working Hardware* as well as the various supported hardware pages on the Bell Labs Plan 9 wiki to help determine if your hardware or VM is supported.


### 3.2 – Known Working Hardware

This list adds to the various supported hardware pages on the Plan 9 from Bell Labs wiki. **Note:** NONE of these lists are all-inclusive. Some drivers have not been tested by 9front developers. The following list consists of hardware we have actually tested.

Some drivers and their options are also documented in `plan9.ini(8)`.

Read: *FQA 1.3.1.2 – New Hardware Support* for information about new hardware drivers in 9front.

#### 3.2.1 – Audio



**AMD FCH Azalia Controller**

vid/did: 1022/780d

**Intel 888 Microsoft UAA bus for HD audio**

vid/did: 8086/284b

**Intel 82801CA/CAM AC97**

vid/did: 8086/2485

**Intel 82801 DB DBM/DA AC 97**

vid/did: 8086/24c5

**Intel 486486 82801IB/IR/IH HD Audio**

vid/did: 8086/293e

**Intel HD NM10/ICH7**

vid/did: 8086/27d8

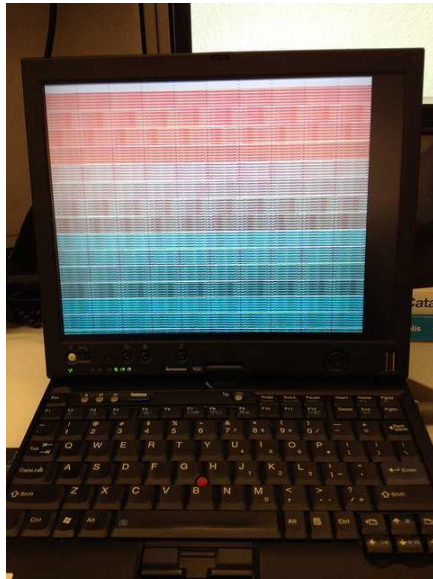
**Intel HD 6 Series/C200 Series**

vid/did: 8086/1c20

**Intel HD 7 Series/C210 Series**

vid/did: 8086/1e20

## 3.2.2 – Graphics



### 3.2.2.1 – AGP

#### **NVidia GeForce FX 5200 128MB VGA output**

```
vid/did: 10de/0322  
monitor=vesa vgasize=1600x1200x32  
monitor=dellst2210 vgasize=1920x1080x32
```

#### **NVidia GeForce FX 5700**

```
vid/did: 10de/0341  
monitor=vesa vgasize=1600x1200x32  
monitor=dellst2210 vgasize=1920x1080x32
```

### 3.2.2.2 – Integrated

#### **ATI Mobility Radeon 7500 128MB DVI/VGA output**

```
vid/did: 1002/4c57  
monitor=vesa vgasize=1024x768x32
```

#### **ATI Mobility Radeon FireGL V3200/X600**

```
vid/did: 1002/3154  
monitor=vesa vgasize=1600x1200x32
```

**ATI RS880**

monitor=vesa vgasize=1280x1024x32

**ATI X1300****Intel Mobile 945GM/GMS/GME, 943/940GML Express**

vid/did: 8086/27a6

monitor=vesa vgasize=1400x1050x32

monitor=x60t vgasize=1400x1050x32

**Intel X3100/GM965/PM965/GL960**

vid/did: 8086/2a03

monitor=vesa vgasize=1680x1050x32

**Intel Mobile Intel 4 Series 4500MHD**

vid/did: 8086/2a42, 8086/2a43

monitor=vesa vgasize=1440x900x32

monitor=x301 vgasize=1440x900x32

**Intel HD 3rd Gen Core processor Graphics Controller**

vid/did: 8086/0166

monitor=vesa vgasize=1366x768x32

monitor=x230 vgasize=1366x768x32

**NVidia GeForce FX Go5200 64M**

vid/did: 10de/0324

monitor=cinema vgasize=1152x768x32

**S3 SuperSavage IX/C 16MB**

vid/did: 5333/8c2e

monitor=t23 vgasize=1024x768x32

monitor=vesa vgasize=1024x768x32

**3.2.2.3 – PCI Express****NVidia GeForce 6200 LE**

vid/did: 10de/0163

monitor=e228wfp vgasize=1680x1050x32

### **NVidia GeForce 8400 GS**

vid/did: 10de/0422  
monitor=vesa vgasize=1680x1050x32

### **NVidia GeForce 8600 GT**

vid/did: 10de/0402  
monitor=vesa vgasize=1600x1200x32

### **NVidia GeForce GTX 550**

vid/did: 10de/0bee  
monitor=vesa vgasize=1600x1200x32

## **3.2.3 – Networking**

### **3.2.3.1 – Ethernet**

#### **3.2.3.1.1 – Integrated**

#### **Broadcom BCM5751M NetXtreme Gigabit**

vid/did: 14e4/167d  
tested 100/1000 mbps

#### **Broadcom BCM5755/5780 NetXtreme Gigabit**

vid/did: 14e4/167b  
tested 100/1000 mbps

#### **Broadcom BCM5782 NetXtreme Gigabit**

vid/did: 14e4/1696

#### **Intel 82540EP Gigabit**

vid/did: 8086/101e  
tested 100/1000 mbps

#### **Intel 82562ET**

tested 10/100 mbps

**Intel 82566MM Gigabit**

vid/did: 8086/1049  
tested 100/1000 mbps

**Intel 82567LM 82567LM-2 Gigabit**

vid/did: 8086/10f5  
tested 100/1000 mbps

**Intel 82573L Gigabit**

vid/did: 8086/109a  
tested: 100/1000 mbps

**Intel 82579LM Gigabit**

vid/did: 8086/1502  
tested: 100/1000 mbps

**Intel 82801CAM PRO/100 VE**

vid/did: 8086/1031  
tested 10/100 mbps

**Realtek RTL8139**

vid/did: 10ec/8139  
tested 10/100/1000 mbps

**Realtek RTL8169/RTL8101E/RTL8102E**

vid/did: 10ec/8136  
tested 10/100/1000 mbps

**3.2.3.1.2 – USB****Beceem Communications CLEAR Stick**

vid/did 198f:8160  
This is a WiMAX device that appears as a USB CDC Ethernet device  
Works with nusb/ether

**RNDIS**

Android phones should work  
Works with nusb/ether

### 3.2.3.2 – WiFi

#### 3.2.3.2.1 – Bridge (external)

##### logear GWU627

802.11n  
connect ethernet port to GWU627

##### Vonets VAP11G

802.11g  
connect ethernet port to VAP11G

#### 3.2.3.2.2 – Mini-PCI

##### Actiontec 800MIP

802.11b  
often branded Lucent WaveLAN  
ether0=type=wavelanpci ssid=MESH station=T42 irq=11

##### Ralink RT2860 802.11b

#### 3.2.3.2.3 – Mini-PCI Express

##### Intel Centrino Advanced-N 6205 Taylor Peak (iwl-6005)

vid/did: 8086/0085  
802.11g  
ether0=type=iwl essid=MESH

##### Intel Centrino Ultimate-N (iwl-6000)

802.11g  
ether0=type=iwl essid=MESH

##### Intel Centrino Wireless-N 100

802.11g  
ether0=type=iwl essid=MESH



### **Intel Centrino Wireless-N 2230**

802.11g  
ether0=type=iwl essid=MESH

### **Intel PRO Wireless 3945ABG (wpi-3945abg)**

802.11g  
ether0=type=wpi essid=MESH

### **Intel WiFi Link 1000/4965/5100/5300 AGN**

802.11g  
ether0=type=iwl essid=MESH

### **Ralink RT3090**

802.11g

### **3.2.3.2.4 – PCI**

### **Ralink RT3090**

802.11b

### **3.2.3.2.5 – PCMCIA**

### **Linksys WPC11**

802.11b  
Prism 2.5  
ISL37300P  
RevA

### **Lucent WaveLAN PC24E-H-FC**

802.11b  
ether0=type=wavelan essid=MESH crypt=off station=x61 irq=11

## **3.2.4 – Tablet Digitizers**

### **3.2.4.1 – Serial**

### 3.2.4.1.1 – Integrated

#### Wacom WACF004

ThinkPad X4\* series tablets

To enable the tablet's serial port in `plan9.ini`:

```
uart2=type=isa port=0x200 irq=5
```

To turn on the tablet:

```
aux/wacom; aux/tablet &
```

#### Wacom WACF008

ThinkPad X6\* series tablets

To enable the tablet's serial port in `plan9.ini`:

```
uart2=type=isa port=0x200 irq=5
```

To turn on the tablet:

```
aux/wacom; aux/tablet &
```

### 3.2.5 – Laptops



#### 3.2.5.1 – Acer

### Aspire 5100 (donated by some poor kid)

cpu: 1795MHz AuthenticAMD AMD-K8, works

ethernet: rtl8139 100Mbps, works

keyboard and touchpad, works

graphics: RS482M ATI RADEON Xpress Series,


monitor=vesa vgasize=1024x768x32 (native resolution not in VESA BIOS); radeon driver untested

wifi: Atheros AR5005G, does not work

audio: SB450 High Definition Audio Controller

### 3.2.5.2 – IBM/Lenovo

#### 3.2.5.2.1 – ThinkPad





## R400

### 7439-1DG

```
cpu: 2527MHz GenuineIntel Core 2/Xeon (cpuid: AX 0x10676 CX 0x8E3FD DX
0xBFEBFBFF)
graphics: Mobile Intel GM45 Express/4500MHD, realemu(8)
monitor=vesa vgasize=1440x900x32
ethernet: Intel 82567LM 82567LM-2 Gigabit (10/100/1000), works
wifi: Intel Wifi Link 5100 AGN Mini-PCI Express, works
disk controller: Intel ICH9M/ME ICH9M/ME AHCI, works
dvd: MATSHITADVD-RAM UJ870A SB04 HE34 068E34 068597, works
audio: Intel 486486 82801IB/IR/IH HD Audio, works
usb: works
mp: mp + sata, ethernet works with *acpi=1
```

## T23

### 2647-HSU

```
cpu: Intel Mobile Pentium III-M 866 MHz, 1.2 GHz
graphics: S3 SuperSavage IX/C 16MB, VGA 1024x768x32, realemu(8)
monitor=vesa vgasize=1024x768x32
ethernet: Intel 82801CAM PRO/100 VE or Intel 82562ET (10/100), works
wifi: Actiontec 800MIP (branded Lucent WaveLAN) Mini-PCI, works
audio: AC97, works
```

## T42

### 2373-BK4

cpu: Intel Pentium M (Dothan) 1.7 GHz  
graphics: ATI Mobility Radeon 7500 32MB, realemu(8)  
monitor=vesa vgasize=1024x768x32  
ethernet: Intel Gigabit Ethernet (10/100/1000), works  
wifi: IBM 11a/b/g Mini-PCI, does not work; replaced with Actiontec 800MIP  
(branded Lucent WaveLAN) Mini-PCI, works  
disk controller: 82801DBM (ICH4-M), IDE DMA works  
audio: AC97, works

## T43p

### 2669-A92

cpu: Intel Pentium M (Dothan) 2.0 GHz (cpuid: AX 0x06D8 CX 0x0180 DX 0xAFE9FBFF)  
graphics: ATI Mobility Radeon FireGL V3200/X600, realemu(8)  
monitor=vesa vgasize=1600x1200x32 with internal LCD  
ethernet: Broadcom BCM5751M (10/100/1000), works  
wifi: replaced with Vonets VAP11G, works  
disk controller: Intel 82801FBM SATA AHCI (ICH6-M): untested  
usb: works  
scram works with \*acpi=1  
mouse button 2 works with trackpad disabled in BIOS

## T60p

### 2007-94U

cpu: Intel Core Duo (Yonah) 2.16GHz (cpuid: AX 0x06E8 CX 0xC1A9 DX 0xBFE9FBFF)  
graphics: ATI MOBILITY FireGL V5200 realemu(8)  
monitor=vesa vgasize=1600x1200x32 with internal LCD  
audio: Intel BA101897 IDT High Definition, untested  
ethernet: Intel 82573L Intel PRO/1000 PL (10/100/1000), works  
wifi: Intel PRO Wireless 3945ABG (wpi-3945abg), may work now with wpi driver  
disk controller: Intel 82801GB/GBM PATA100, 82801GBM/GHM AHCI, works  
usb: works  
mp: mp + sata, ethernet, usb works with \*acpi=  
mouse button 2 works with trackpad disabled in BIOS

## T61

### 7659-CTO

cpu: Intel Core 2 Duo (Merom) 2.0 GHz FSB, 2-4MB L2 Cache CPU  
graphics: Intel GMA X3100, realemu(8)  
monitor=vesa vgasize=1280x800x32 with internal LCD;  
monitor=vesa vgasize=1680x1050x32 with VGA or Mini Doc DVI output  
and external monitor

ethernet: Intel 82566MM (10/100/1000), works  
wifi: Intel Wireless WiFi Link 4965 AGN Mini-PCI Express, should work with iwlmwmi driver  
disk controller: Intel 82801HBM/HEM PATA, Intel 82801HBM SATA AHCI (ICH8-M): IDE DMA works, SATA works  
usb: works  
mp: mp + sata, ethernet, usb works with \*acpi=  
mouse button 2 works with trackpad disabled in BIOS

#### 7661-12U

cpu: Intel Core 2 Duo (Merom) 2.0 GHz FSB, 2-4MB L2 Cache CPU  
graphics: Intel GMA X3100, realemu(8)  
monitor=vesa vgasize=1280x800x32 with internal LCD;  
monitor=vesa vgasize=1680x1050x32 with VGA or Mini Doc DVI output and external monitor  
ethernet: Intel 82566MM (10/100/1000), works  
wifi: Intel PRO Wireless 3945ABG (wpi-3945abg), may work now with wpi driver; replaced with Wavelan PC24E-H-FC PCMCIA, works  
disk controller: Intel 82801HBM/HEM PATA, Intel 82801HBM SATA AHCI (ICH8-M): IDE DMA works, SATA works  
usb: works  
mp: mp + sata, ethernet, usb works with \*acpi=  
mouse button 2 works with trackpad disabled in BIOS

#### T400

#### 6475-EC7

cpu: 2261MHz GenuineIntel Core 2/Xeon (cpuid: AX 0x10676 CX 0x8E3FD DX 0xBFEBFBFF)  
graphics: Intel Corporation Mobile 4 Series, realemu(8)  
monitor=vesa vgasize=1440x900x32 with internal LCD  
ethernet: i82567: 1000Mbps, works  
wifi: Intel WiFi Link 5100 AGN Mini-PCI Express, works  
disk controller: Intel ICH9M/ME AHCI, works  
audio: Intel HDA, should work

#### T410i

#### 2518-4QG

cpu: Intel(R) Core(TM) i5 CPU M 430 @ 2.27GHz  
graphics: Intel Graphics Media Accelerator HD, realemu(8)  
monitor=vesa vgasize=1280x800x32 with internal LCD  
ethernet: Intel 82577LM Gigabit, works  
wifi: unknown, reportedly works  
audio: unknown, reportedly works

## T420s

### 4171-53U

cpu: Intel® Core™ i5-2540M (2.6GHz, 3MB L3, 1333MHz FSB) (cpuid: AX 000206A7 CX 17BAE3FF DX BFEBFBFF)  
graphics: Intel HD Graphics 3000 (integrated Sandy Bridge GPU), realem(8)  
monitor=vesa vgasize=1600x900x32  
ethernet: Intel 82579LM Gigabit, works  
audio: Intel HD 6 Series/C200 Series, works  
wifi: Intel Centrino Advanced-N 6205 Taylor Peak, etheriwl (firmware: iwn-6005), works



## X41 tablet

cpu: Intel Pentium M (Dothan) 1.6GHz LV (778) L2 2 MB cache  
graphics: Intel GMA900, realem(8)  
monitor=vesa vgasize=1024x768x32  
tablet: WACF004, works  
ethernet: BCM5751M (10/100/1000), works  
wifi: Intel PRO/Wireless 2915ABG Mini-PCI, does not work

## X60s


### 1704-GL5

coreboot  
cpu cpu0: 1663MHz GenuineIntel P6 (cpuid: AX 0x06E8 CX 0xC1A9 DX 0xBFE9FBFF)  
graphics: Intel 945GM, realem(8)  
monitor=vesa vgasize=1024x768x16  
ethernet: Intel 82573L Intel PRO/1000 PL, works  
audio: Intel HDA, untested  
wifi: Ralink RT3090, works

## X60 Tablet

### 6363-CTO

cpu: Intel Core Duo (Yonah) L2400 LV 1.66 GHz (2MB Cache) cpu0: 1663MHz GenuineIntel P6 (AX 000006EC CX 0000C1A9 DX BFE9FBFF)  
lcd: replaced with HV121P01-100 (1400x1050)  
graphics: Intel Graphics Media Accelerator 950, realemu(8)  
monitor=x60t vgasize=1400x1050x32;  
monitor=vesa vgasize=1280x1024x32 (native 1400x1050 resolution not in VESA BIOS)  
tablet: WACF008, works  
ethernet: Intel 82573L (10/100/1000), works  
audio: Intel HD Audio with AD1981HD codec, speaker and green lineout work with `echo pin 5,3 >/dev/audioctl`  
wifi: Intel WiFi Link 5100 AGN, works — flashed with custom BIOS to remove WiFi card whitelist  
disk controller: Intel 82801GBM/GHM (ICH7-M Family) SATA Controller [AHCI mode], untested  
mp: mp + sata, ethernet works with `*acpi`



## X61s


cpu: Intel Core 2 Duo  
graphics: Intel GM965/GL960, realemu(8)  
monitor=vesa vgasize=1024x768x32  
ethernet: Intel 82566MM (10/100/1000), works  
wifi: Intel PRO/Wireless 4965 AG or AGN '[Kedron]' Mini-PCI Express, untested, should work with `iwl` driver  
disk controller: Intel 82801HBM SATA (ICH8-M): IDE DMA works, SATA works  
mp: mp + sata, ethernet works with `*acpi=`



## X61 Tablet

### 7767-01U

cpu: Intel Core 2 Duo CPU L7700 (1.80 GHz)  
lcd: replaced with HV121P01-100 (1400x1050)  
graphics: Intel GMA X3100, realemu(8)  
monitor=x60t vgasize=1400x1050x32;  
monitor=vesa vgasize=1280x1024x32 (native 1400x1050 resolution not  
in VESA BIOS)  
tablet: WACF008, works  
ethernet: Intel 82566MM (10/100/1000), works  
audio: Intel HDA, works  
wifi: Intel Centrino Advanced-N 6205 Taylor Peak (iwl-6005), works — flashed  
with custom BIOS to remove WiFi card whitelist  
disk controller: Intel 82801HBM SATA (ICH8-M): IDE DMA works, SATA works mp:  
mp + sata, ethernet works with \*acpi=



## X131e

### 3368-2FU

cpu: 1397MHz GenuineIntel P6 (cpuid: AX 0x206A7 CX 0x15BAE3BF DX 0xBFEBF-  
BFF)  
graphics: Intel HD Graphics 3000, realemu(8)  
monitor=vesa vgasize=1366x768x32  
ethernet: Realtek RTL8167 PCIe Gigabit Ethernet: works  
audio: Intel HDA, works

## X200

### 7455-A54

cpu: Intel Core 2 Duo  
graphics: Mobile Intel GMA 4500MHD, realem(8)  
monitor=vesa vgasize=1280x800x32  
ethernet: Intel 82567LF Gigabit (10/100/1000), works  
wifi: Intel WiFi Link 5150, works  
disk controller: Intel ICH9M/ME ICH9M/ME AHCI, works  
audio: Intel 486486 82801IB/IR/IH HD Audio, works  
usb: works

## X200s

### 7466-3SG

cpu: 1862MHz GenuineIntel Core 2/Xeon (cpuid: AX 0x10676 CX 0x8E3FD DX 0xBFEBFBFF)  
graphics: Mobile Intel GM45 Express/4500, realem(8)  
monitor=vesa vgasize=1280x800x32  
ethernet: Intel 82567LM 82567LM-2 Gigabit (10/100/1000), works  
wifi: Intel WiFi Link 5300 AGN Mini-PCI Express, works  
disk controller: Intel ICH9M/ME ICH9M/ME AHCI, works  
audio: Intel 486486 82801IB/IR/IH HD Audio, works  
usb: works  
mp: mp + sata, ethernet works with \*acpi=1

## X201

### 3323-DBG


cpu: 2661MHz GenuineIntel P6  
graphics: Mobile Intel GMA 5700MHD, realem(8)  
monitor=vesa vgasize=1280x800x32  
ethernet: Intel Corporation 82577LM Intel 82577LM Gigabit, works  
wifi: Intel Centrino Ultimate-N 6300 AGN, works  
disk controller: Intel Corporation PCH (Ibex Peak) SATA AHCI, works  
audio: Intel HDA, works  
ssd: INTEL SSDSC2BW180A3L, works  
usb: works

## X220

### 4291-4CG

cpu: 2791MHz GenuineIntel P6 (cpuid: AX 0x206A7 CX 0x17BAE3FF DX 0xBFEBFBFF)  
graphics: Intel HD 3000, realem(8)  
monitor=vesa vgasize=1366x768x32  
ethernet: Intel 82579 (10/100/1000), works  
audio: Intel HDA, works  
wifi: Intel Centrino Advanced-N 6205, works


disk controller: works  
ssd: INTEL SSDSA2BW160G3, works  
usb: works



X230

2306-CTO

cpu: Intel Core i5-3320M (2.60 GHz, 3MB L3, 1600MHz FSB), cpuid: AX 0x306A9  
CX 0x77BAE3FF DX 0xBFEBFBFF  
graphics: Intel HD 3rd Gen Core processor Graphics Controller, realemu(8)  
monitor=x230 vgasize=1366x768x32;  
monitor=vesa vgasize=1366x768x32  
ethernet: Intel 82579LM Gigabit (10/100/1000), works  
wifi: Intel Centrino Advanced-N 6205 Taylor Peak, etheriwl (firmware: iwn-6005),  
works  
disk controller: Intel 7 Series Chipset Family 6-port SATA Controller AHCI mode,  
ahci, works  
usb: Intel 7 Series/C210 Series Chipset Family USB Enhanced Host Controller #1,  
ehci, works  
audio: Intel 7 Series/C210 Series Chipset Family High Definition Audio Controller,  
works  
efi: works



## X230 Tablet

### 3434-CTO

cpu: Intel Core i5-3320M (Ivy Bridge), 2.6 GHz, 3 MB Shared L3 Cache, 2C/4T, 35 W TDP, 22 nm (cpuid: AX 0x306A9 CX 0x77BAE3FF DX 0xBFEBFBFF)  
graphics: Intel HD 4000, realemu(8)  
monitor=vesa vgasize=1366x768x32  
tablet: USB, untested  
ethernet: Intel 82579 (10/100/1000), works  
audio: Intel HDA, works  
wifi: Intel 802.11b/g/n, might work  
disk controller: works  
ssd: SAMSUNG SSD 830 Series CXM03B1Q S0XYNEAC774074 128GB, works  
usb: disable USB3 in BIOS, works


## X301

### 2776-P4U

cpu: Intel Core 2 Duo SU9400 1.4GHz 3MB cache, cpuid: AX 0x1067A CX 0x408E3FD DX 0xBFEBFBFF  
graphics: Mobile Intel GM45 Express/4500MHD, realemu(8)  
monitor=x301 vgasize=1440x900x32;  
monitor=vesa vgasize=1440x900x32  
ethernet: Intel 82567LM 82567LM-2 Gigabit (10/100/1000), works  
wifi: Intel WiFi Link 5100 AGN Mini-PCI Express, works  
disk controller: Intel ICH9M/ME ICH9M/ME AHCI, works  
ssd: Samsung MMCRE64G8MPP-0VA 64GB, works  
dvd: Matsushita DVD-RAM UJ-844, works  
audio: Intel HD 486486 82801IB/IR/IH, works  
usb: works  
mp: mp + sata, ethernet works with \*acpi=1

## 2776-P6U

cpu: Intel Core 2 Duo SU9400 1.4GHz 3MB cache, cpuid: AX 00010676 CX 0008E3FD DX BFEBFBFF  
graphics: Mobile Intel GM45 Express/4500MHD, realemu(8)  
monitor=x301 vgasize=1440x900x32;  
monitor=vesa vgasize=1440x900x32  
ethernet: Intel 82567LM Gigabit (10/100/1000), works  
wifi: Intel WiFi Link 5100 AGN Mini-PCI Express, works  
disk controller: 82801IBM/IEM (ICH9M/ICH9M-E) 4 port SATA Controller AHCI mode, works  
audio: Intel HD 82801I (ICH9 Family), works  
usb: works  
mp: mp + sata, ethernet works with \*acpi=1



## 3.2.5.3 – Toshiba

### 3.2.5.3.1 – Satellite

## M30-S309

cpu: 1397MHz GenuineIntel P6 (AX 00000695 CX 00000180 DX A7E9F9BF)  
graphics: NVidia GeForce FX Go5200 64M,  
monitor=cinema vgasize=1152x768x32  
ethernet: Intel 82801DB PRO/100 VE (MOB) (i82557), works  
disk controller: Intel 82801DBM (ICH4-M) IDE Controller, works  
audio: Intel 82801DB/DBL/DBM (ICH4/ICH4-L/ICH4-M) AC'97 Audio Controller, works  
usb: Intel 82801DB/DBL/DBM (ICH4/ICH4-L/ICH4-M) USB UHCI Controller, untested

## 3.2.6 – Desktops



### 3.2.6.1 – eMachines

T3302 cpu: AMD Sempron 3300+ 2GHz

chipset: VIA K8M800

graphics: VIA S3 UniChrome, replaced with NVidia GeoForce FX 5700 128MB DVI output, 1920x1080x32

ethernet: 3Com 3C905-TX Fast Etherlink 10/100 PCI TX

audio:

usb: works

### 3.2.6.2 – Igel

4210 LX Winestra

### 3.2.6.3 – Soekris

net6501-70

cpu: Intel Atom E680 1.6Ghz, both pc and pc64 work

ethernet: 4x Intel 82574L Gigabit Ethernet, works

usb: works

serial console: works, use `console=0 b19200` in `plan9.ini`. 9boot hangs without a serial cable attached; disable `uartputc` as a workaround

this machine does not have ACPI

### 3.2.6.4 – IBM/Lenovo

#### 3.2.6.4.1 – ThinkCentre

M55

#### 8810–D3U

cpu: Intel Core 2 Duo  
graphics: Intel GMA 3000 internal (untested), replaced with NVidia GeForce 8400GS  
DVI output, realem(8) monitor=vesa vgasize=1680x1050x32  
ethernet: Broadcom BCM5755 (10/100/1000), works  
audio: Intel HDA, works  
usb: works  
mp: mp + sata, ethernet, usb works with \*acpi=

### 3.3 – Virtual Machines

9front has been tested on several virtual machines. Details below.

**Note:** As a general rule it is a good idea to manually specify a unique MAC address for each virtual machine instance running on the network, to avoid collisions.

#### 3.3.1 – Qemu

The following is tested with qemu 1.5.0 and 2.0.50 running on Linux, using *FQA 3.3.3 – virtio* for disk and network.

##### 3.3.1.1 – Installation

Create a sparse disk image:

```
qemu-img create -f qcow2 9front.qcow2.img 30G
```

Boot the 9front.iso:

```
qemu-system-x86_64 -cpu host -enable-kvm -m 1024 \  
-net nic,model=virtio,macaddr=52:54:00:00:EE:03 -net user \  
-device virtio-scsi-pci,id=scsi \  
-drive if=none,id=vd0,file=9front.qcow2.img -device scsi-hd,drive=vd0 \  
-drive if=none,id=vd1,file=9front.iso -device scsi-cd,drive=vd1,bootindex=0
```

Finally, see: *FQA 4.3 – Performing a simple install*

##### 3.3.1.2 – Post-Installation Booting

```
qemu-system-x86_64 -cpu host -enable-kvm -m 1024 \  
-net nic,model=virtio,macaddr=52:54:00:00:EE:03 -net user \  
-device virtio-scsi-pci,id=scsi \  
-drive if=none,id=vd0,file=9front.qcow2.img -device scsi-hd,drive=vd0
```

### 3.3.1.2.1 – Multiboot

Multiboot can be used to start the 9front kernel directly, skipping the bootloader step:

```
-qemu -kernel 9pcf -initrd plan9.ini
```

### 3.3.1.4 – Networking

User networking is the default and works the same on every platform. More advanced options are particular to specific host operating systems; several are described below.

**Note:** On many operating systems ICMP is limited to the superuser. One consequence is that a VM running with guest networking cannot ping remote hosts.

#### 3.3.1.4.1 – Linux VDE

Install vde2.

Setup a tap interface:

```
sudo tuncctl -u $USER -t tap0
```

Start a virtual switch connected to the tap interface:

```
vde_switch --tap tap0 -daemon
```

Connect the switch to the network of the host. Use DHCP:

```
slirpvde --dhcp --daemon
```

When booting 9front, add the following to the `qemu` command line arguments:

```
-net vde
```

#### 3.3.1.4.2 – OpenBSD tun(4)

Tested: OpenBSD/amd64 4.9-CURRENT, qemu-0.9.1p18, kqemu-1.3.0pre11p3

Configure a tun(4) interface with an IP address selected for your `qemu` network. Configure bridging between the tun(4) interface and your machine's external interface. Finally, configure NAT or rules to pass traffic for the `qemu` network in `pf.conf`. Reboot.

To launch `qemu` as a regular user:

```
sudo sh -c "sudo -C 4 -u $USER qemu -m 768 \  
-net nic,vlan=0,model=rtl8139,macaddr=52:54:00:12:33:36 \  
-net tap,vlan=0,fd=3,script=no -no-fd-bootchk \  
-hda 9front.qcow2.img 3<>/dev/tun0"
```

Read: `hostname.if(5)`, `pf.conf(5)`, `tun(4)`



### 3.3.1.4.3 – Windows TAP

This is tested with the qemu for windows distribution. Download and run the installer from openvpn to install the windows TAP driver. Create a new TAP interface with the "Add a new TAP virtual ethernet adapter" from the openvpn start menu. Go to the network manager and rename that new TAP interface to something more sane like: "qemu-tap". Configure ip addresses or bridge that interface with the network manager.

Now you should be able to run qemu on that interface:

```
qemu.exe -net nic -net tap,ifname="tap-qemu" ...
```

### 3.3.1.5 – Audio

Run qemu with the flag `-soundhw sb16` and put the following line in `plan9.ini`:

```
audio0=type=sb16 port=0x220 irq=5 dma=5
```

**Note:** `irq` and `dma` values may vary.

### 3.3.1.6 – Graphics

Use `monitor=vesa`

**Note:** Some versions of QEMU running on OSX have exhibited graphical glitches when using a 16-bit color mode (for example: 1024x768x16. Try a 32-bit mode instead (for example: 1024x768x32).

### 3.3.2 – Virtualbox

Don't use Virtualbox. It tends to break between versions.



**3.3.2.1 – Ethernet** The emulated "Intel PRO/1000 MT Server" ethernet controller is known to work.

### 3.3.2.2 – Audio

Put the following in `plan9.ini`:

```
audio0=type=sb16
```

**3.3.2.3 – Graphics** Use `monitor=vesa`

### 3.3.2.4 – Known Working Versions

4.3.14 r95030 on Windows 7

4.3.16 on Mac OS X

4.3.18 r96516 on Linux x86\_64 kernel 3.14.22

4.3.18 on Windows 7:

just tried with vbox 4.3.18 on windows7. 9front boots fine in BIOS mode, but the PCnet nic dosnt work. reason is that vbox pllx pci irq routing is fucked so the ethernet doesnt get interrupts. if i boot with `*nopcirouting=1`, it works fine. theres a option to select the chipset so i tried ICH9 with IO-APIC enabled. normal mp mode fails because of broken mp tables, but works with `*acpi=`. also, it works with UEFI mode (which always uses ACPI). the usual intel mt server nic also works (thats what is usually recommended for working around the broken ethernet).

pci routing issue has been fixed in latest kernel, should be available in iso release after 3960.

#### 4.3.20 r96996 on Mac OS X 10.6.8/10.9 and Ubuntu 14.04/14.10:

General -> Basic Type: Other Version: Other/Uknown

System -> Motherboard Chipset: PIIX3 Pointing Device: PS/2 Mouse Extended Features:  Enable I/O APIC

System -> Processor Extended Features:  PAE/NX (not sure this matters)

System -> Acceleration  Enable VT-x/AMD-V  Enable Nested Paging

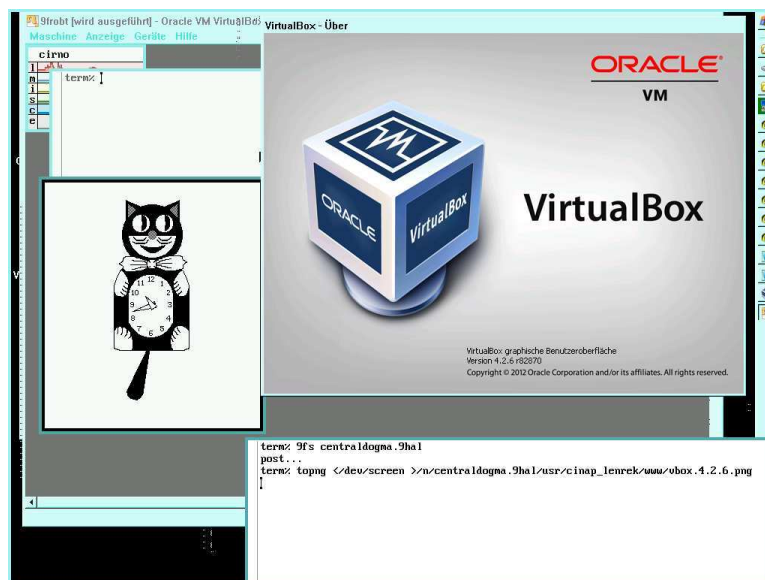
Display -> Video Extended Features:  Enable 3D Acceleration (not sure this matters)

Storage -> Attributes Name: IDE Type: PIIX4  Use Host I/O Cache

Audio ->  Enable Audio Host Audio Driver: CoreAudio (Can be PulseAudio or otherwise for Linux, etc. Shouldn't be hard to set this) Audio Controller: Soundblaster 16

Network -> Adapter 1 Attached to: NAT -> Advanced Adapter Type: Intel PRO/1000 MT Server Promiscuous Mode: Deny (Not sure this matters)

**Note:** Enabling USB 2.0 Controll in 'Ports -> USB' works just fine in 9front, mounting under /shx flawlessly as long as the host has the Virtualbox Extension Pack running.




### 3.3.3 – Virtio

Current versions of qemu/kvm and virtualbox as of 3.1 support faster paravirtualized devices. Presently, 9front provides drivers for virtio hard disk and network.

The virtio-blk disk device should show up as: `/dev/sdF0`

The virtio-scsi disk device should show up as: `/dev/sd00`

## FQA 4 – 9front Installation Guide



**ACHTUNG!** This document (and the 9front installer) pertains only to the 386 install target. For other architectures, read: 8.9 – Bootstrapping architectures not included on the ISO



## 4.1 – Pre-installation Checklist

Before you start your install, you should have some idea what you want to end up with. You will want to know the following items, at least:

**Machine name**, hereafter referred to as `sysname`.

**Hardware installed and available.** Read: *FQA 3.3 – Selecting Hardware*. Check *FQA 3.2 – Known Working Hardware* as well as the various supported hardware pages on the Bell Labs Plan 9 wiki to help determine if your hardware or VM is supported.

**Network settings, if not using DHCP:** Gather the requisite information for your network (IP, netmask, default gateway, etc.). If you plan to install in a VM, see *FQA 3.3 – Virtual Machines*.

In case of trouble, see: *FQA 9 – Troubleshooting*, *FQA 2.2.1 – Mailing List Archives*

## 4.2 – Creating bootable 9front install media

### 4.2.1 – ISO image

Obtain the `9front.iso`.

Read: *1.11.1 – Mirrors*

## 4.2.2 – USB drive

Booting from USB disk is handled the same way as booting from internal hard drive. In general, the boot process is able to use an ISO image stored on a FAT file system as its root file system, so all that is needed is to install the `9boot(8)` bootloader, a kernel and the `9front.iso` on the USB disk.

### 4.2.2.1 – Creating on Plan 9

The path to your USB device will look something like this: `/dev/sdUxxxxx` where `xxxxx` is the unique name of your usb device.

Build and install the mbr and boot loader:

```
cd /sys/src/boot/pc
mk 9bootfat mbr pbs
disk/mbr -m mbr /dev/sdUxxxxx/data
```

If it doesn't already exist, create and format a FAT partition. Don't forget to set the FAT partition active in the `fdisk` menu:

```
disk/fdisk -b /dev/sdUxxxxx/data
disk/format -b pbs -d -r 2 /dev/sdUxxxxx/dos
```

Mount the USB device:

```
dosrv -f /dev/sdUxxxxx/dos sdos
mount -c /srv/sdos /n/dos
```

Create a suitable `/n/dos/plan9.ini`:

```
bootfile=9pcf
mouseport=ask
monitor=ask
vgasize=ask
```

**Note:** By convention, Plan 9 usually expects text files to end with a newline. If the last line of `plan9.ini` is not a newline, it could fail to be parsed correctly at boot time.

Copy files to the USB device:

```
cp /386/9bootfat /n/dos
chmod +al /n/dos/9bootfat # defrag magic
cp /386/9pcf /n/dos
cp /path/to/9front.iso /n/dos
```

Optional for EFI systems:

```
cd /sys/src/boot/efi; mk install
mkdir -p /n/dos/efi/boot
cp /386/boot*.efi /n/dos/efi/boot
```

Unmount the USB device:

```
umount /n/dos
rm -f /srv/sdos
```

Boot the device.

Read: `prep(8)`

#### 4.2.2.2 – Creating on Linux

**Note:** There are numerous methods for creating bootable USB devices on Linux. This is one.

Obtain `mbr`, `pbs`, `9bootfat` and `9pcf` binaries. Either copy them from the distributed ISO image (the files are located under `/386/`) or build up to date binaries on a 9front system. Read: *FQA 5.2.2 – Building from source*

On the Linux system, create a `$dir` with that contains the following files:

```
9front.iso
9pcf
plan9.ini (as described above)
9bootfat
mbr
pbs
```

Optional for EFI systems, create the directory `boot` in `$dir` and copy the files `bootia32.efi` and `bootx64.efi` into it.

Use the `makebootfat` tool to create bootable device (ie. `/dev/sdc`):

```
$ makebootfat -m $dir/mbr -b $dir/pbs -o /dev/sdc $dir
```

Boot the device.

#### 4.2.2.3 – Bootargs

At the `bootargs` prompt you'll need to enter the path to the ISO on the USB device. It will look something like this:

```
bootargs=local!/shr/sdUxxxxx/9front.iso
```

If a USB device is not listed, escape to a shell with `!rc`, and `ls /shr` to find it.

Read: 9.5.1 – Devices not recognized or not working

### 4.3 – Performing a simple install

There is nothing magical about installing Plan 9. It is simply a matter of populating a Plan 9 file system (`cwfs` or `hjfs`) and arranging a bootstrap to eventually load a Plan 9 kernel that can then use that file system as its root.

In most cases, the file server is the only machine that needs to have a disk. Once the



initial file server is running, setting up an auth server, and enabling 'bootp' and 'tftp' for PXE booting, will allow Plan 9 terminals and cpu servers to load kernels from the file server and share its file system over the network.

**Note:** This guide describes the default installation of a terminal with disk, which is an amalgamation of a normal Plan 9 network, but is sufficient for exploring the basics of using Plan 9. Configuration of additional services (such as a file server with networking listeners, a cpu server, etc.) is outlined in *FQA 7 – System Management*.

To install 9front, boot the 9front.iso image and follow the steps below.

#### 4.3.1 – boot

Successfully booting the system will result in basic information about the state of the system being printed on screen:

```
Plan 9
126 holes free
00018000 0009f000 552960
00485000 0cf2c000 212496384
213049344 bytes free
cpu0: 3395MHz GenuineIntel P6 (cpuid: AX 0x206A7 CX 0x82982203 DX 0xFEBFBFF)
ELCR: 0E20
#10: AMD79C970: 10Mbps port 0x2000 irq 10: 000c291d0baf
mylex ctlr @ port 0x10c0: 32-bit wide SCSI host adapter
512M memory: 207M kernel data, 304M user, 929M swap
nusb/usbd: /dev/usb: no hubs
```

#### 4.3.2 – bootargs

A list of attached storage devices is printed, followed by a prompt asking for a Plan 9 partition to boot from:

```
/dev/sdC0: VMware Virtual IDE Hard Drive
/dev/sdC0/data
/dev/sdD0: VMware Virtual IDE CDROM Drive
/dev/sdD0/data 9660
bootargs is (tcp, il, local!device) [local!/dev/sdD0/data]
```

In most cases, the suggested default will correspond to the device used to boot the ISO image.

#### 4.3.3 – user

Next, a prompt asks for a username:

```
user[glenda]:
```

Glenda is the default `hostowner` name. Don't change this until you know what you are doing. Hit `enter` to accept the default.

**Note:** The `hostowner` differs from the concept of `root` on a UNIX system, where a single user `root` may take control of all processes *and* files on the system. By contrast,

even the `hostowner` of a Plan 9 file server cannot violate file permissions on the file system, except when permissions checking is disabled on the console or when entering special commands at the console of the file server. The `hostowner` controls only the *processes* running on the local machine (in the case of the `filserver`, the file server process itself is obviously owned by the system's `hostowner` ). This fundamental separation between control of processes and file permissions is exploited throughout the system, but can be confusing for users coming from a UNIX background.

#### 4.3.4 – `vgasize`, `monitor`, `mouseport`

The next set of prompts deal with graphics display and the mouse:

```
vgasize is (text, 640x480x8, 1024x768x16, ...) [1024x768x16]
monitor is (vesa, xga, lcd, ...) [vesa]
mouseport is (ps2, ps2intellimouse, 0, 1, 2) [ps2]
```

The boot process prompts for the environment variables `$vgasize`, `$monitor` and `$mouseport`, and the installer will later write those values to the system's `plan9.ini`, from which they are loaded on subsequent system bootup.

Setting `$monitor` to anything besides `vesa` will bypass the emulated VESA BIOS and attempt to use a native VGA driver for the video card in question. Read: `/lib/vgadb` for a list of monitors and video cards that are already known by the system; and the man pages `vga(3)`, `vga(8)` and `vgadb(6)` for more information about how graphical displays are configured.

The default `$mouseport` of `ps2` is sufficient in most cases. Set it to `ps2intellimouse` on laptops or for mice that have a scrollwheel.

**Note:** Some laptops require the trackpad to be disabled in the BIOS in order for mouse button 2 (the center button) to function in Plan 9. In a pinch, mouse button 2 may be simulated by holding down the `shift` key while clicking mouse button 3 (the right button).

After the install, changes intended to persist across reboots should be added to `plan9.ini`.

Examples:

Values as they appear in `plan9.ini`:

```
monitor=vesa
vgasize=1024x768x16
mouseport=ps2intellimouse
```

Read: `mouse(8)`, `vga(3)`, `vgadb(6)`, `plan9.ini(8)`, `realemu(8)`, `vga(8)`, *FQA 7.2.2 – How do I modify plan9.ini?*

After the `mouseport` prompt is answered, the boot process will attempt to start the Plan 9 graphical environment, `rio(1)`, opening a `stats(8)` window and a `rio` window on top of a gray desktop background.



#### 4.3.4.1 – Changing screen resolution

At this point it may be desirable to change the screen resolution. To change video mode from the command line:

vesa:

```
@{rfork n; aux/realemu; aux/vga -p} # obtain a list of vesa bios modes  
@{rfork n; aux/realemu; aux/vga -m vesa -l 1024x768x16}
```

vga:

```
aux/vga -m dellst2210 -l 1920x1080x32
```

#### 4.3.5 – inst/start

Installation is performed by the rc scripts in `/rc/bin/inst`

**Note:** Any task may be repeated by manually entering its name at the next `Task to do` prompt.

```
term% inst/start  
Tue Jul 17 12:38:50 CET 2012 Installation process started
```

The following `Task to do` steps are handled one at a time:

### 4.3.6 – configfs

You can install the following types of systems:

```
    cwfs64x      the cached-worm file server
    hjfs         the new 9front file server (experimental!)
```

File system (cwfs64x, hjfs)[cwfs64x]:

**Note:** If you are installing to a disk of less than 12GB, you should choose `hjfs` (its cache and permanent storage are both located on the same partition, and so require less overall disk space).

Read: *FQA 3.3 – What is an appropriate first system to learn 9front on?*

In this example we will press `enter` to accept `cwfs64x` as the default.

### 4.3.7 – partdisk

The following disk devices were found.

```
sdC0 – VMware Virtual IDE Hard Drive
      empty                0 3916          (3916 cylinders, 29.99 GB)
```

```
sdD0 – VMware Virtual IDE CDROM Drive
```

Disk to partition (sdC0, sdD0)[no default]:

Enter the media you wish to install to.

```
Disk to partition (sdC0, sdD0)[no default]: sdC0
The disk you selected HAS NO master boot record on its first sector.
(Perhaps it is a completely blank disk.)
You need a master boot record to use the disk.
Should we install a default master boot record?
```

Assuming a blank disk image, install a fresh mbr:

```
Install mbr (y, n)[no default]: y
```

```
This is disk/fdisk; use it to create a Plan 9 partition.
If there is enough room, a Plan 9 partition will be
suggested; you can probably just type 'w' and then 'q'.
```

```
cylinder = 8225280 bytes
>>>
```

For this example we will use the entire disk. Accept the defaults.

```
>>> w
>>> q
```

#### 4.3.8 – prepdisk

The following Plan 9 disk partitions were found.

```
/dev/sdC0/plan9
  empty                0 62910477    (62910477 sectors, 29.99 GB)
```

Plan 9 partition to subdivide (/dev/sdC0/plan9)[/dev/sdC0/plan9]:

Use the Plan 9 partition created in the previous step. Hit 'enter' to select the default.

This is disk/prep; use it to subdivide the Plan 9 partition.  
If it is not yet subdivided, a sensible layout will be suggested;  
you can probably just type 'w' and then 'q'.

```
no plan9 partition table found
9fat 204800
nvram 1
other 8957953
fscache 8957953
fsworm 44789770
>>>
```

Again, accept the defaults.

```
>>> w
>>> q
```

#### 4.3.9 – mountfs

The please choose your cwfs64x partitions

Hit enter to select each partition at the prompt.

```
--rw-r----- S 0 glenda glenda 4586471936 Jul  4 13:28 /dev/sdC0/fscache

Cwfs cache partition (/dev/sdC0/fscache)[/dev/sdC0/fscache]:
--rw-r----- S 0 glenda glenda 22932362240 Jul  4 13:28 /dev/sdC0/fsworm

Cwfs worm partition (/dev/sdC0/fsworm)[/dev/sdC0/fsworm]:
--rw-r----- S 0 glenda glenda 4586471936 Jul  4 13:28 /dev/sdC0/other

Cwfs other partition (/dev/sdC0/other)[/dev/sdC0/other]:
```

Since this is a fresh install, we choose yes to ream (format) the file system:

```
Ream the file system? (yes, no)[no]: yes
Starting cwfs64x file server for /dev/sdC0/fscache
Reaming file system
bad nvram key
bad authentication id
bad authentication domain
nvrcheck: can't read nvram
config: config: config: auth is now disabled
config: config: config: config: config: config: current fs is "main"
cmd_users: cannot access /adm/users
63-bit cwfs as of Wed Jul  4 00:59:30 2012
      last boot Tue Jul 17 13:34:57 2012
Configuering cwfs64x file server for /dev/sdC0/fscache
Mounting cwfs64x file server for /dev/sdC0/fscache
% mount -c /srv/cwfs /n/newfs
Mounting cwfs64x file server for /dev/sdC0/other
% mount -c /srv/cwfs /n/other other
```

#### 4.3.10 – configdist

Are you going to download the distribution  
from the internet or do you have it on local media?

Distribution is from (local, net)[local]:

We are installing from local media, so, local. Hit 'enter' to accept the default.

**Note:** The net option is somewhat experimental, and may work. Attempt it at your own risk.

#### 4.3.11 – confignet

You can connect to the internet via  
a local ethernet or a dial-up PPP connection.

Interface to use (ether, ppp)[ether]:

Please choose a method for configuring your ethernet connection.

manual – specify IP address, network mask, gateway IP address  
dhcp – use DHCP to automatically configure

Configuration method (manual, dhcp)[dhcp]:

##### 4.3.11.1 – dhcp

Hit enter to move on to the next task.

##### 4.3.11.2 – manual

If you chose manual, enter values that are appropriate for your network.

```
Configuration method (manual, dhcp)[dhcp]: manual
ip address [no default]: 10.0.2.15
network mask [no default]: 255.255.255.0
gateway address [no default]: 10.0.2.2
```

#### 4.3.12 – mountdist

```
Please wait... Scanning storage devices...
/dev/sdC0/9fat
/dev/sdC0/data
/dev/sdC0/fscache
/dev/sdC0/fsworm
/dev/sdC0/other
/dev/sdD0/data
```

The following storage media were detected.  
Choose the one containing the distribution.

```
/dev/sdD0/data (iso9660 cdrom)
```

```
Distribution disk (/dev/sdD0/data, /dev/sdC0/fscache, /)[/]:
```

The CD-ROM is already mounted at /, so we hit enter to choose the default.

```
% mount /srv/boot /n/distmedia
```

```
Which directory contains the distribution?
Any of the following will suffice (in order of preference):
- the root directory of the cd image
- the directory containing 9front.iso
- the directory containing 9front.iso.bz2
```

```
Location of archives [/]:
```

And again, the root directory of the CD-ROM is already mounted at /, so hit enter to choose the default.

```
% mount /srv/boot /n/distmedia
```

```
Which directory contains the distribution?
Any of the following will suffice (in order of preference):
- the root directory of the cd image
- the directory containing 9front.iso
- the directory containing 9front.iso.bz2
```

```
Location of archives [/]:
```

Hit enter.

**4.3.13 – copydist** Hit enter at the `copydist` prompt to begin the process of copying the distribution files from the install media to the hard disk.

```
processing /sys/lib/sysconfig/proto/allproto
```



**ACHTUNG!** Currently, there is no progress meter. For systems without DMA, the `copydist` task may exceed one hour in duration. Disk activity may be verified by inspecting the `stats(8)` window.



Eventually, you should see the following:

```
file system made
```



This indicates that the system files have completed copying to the install target.

#### 4.3.14 – sysname

Setup network configuration

sysname [cirno]:

Use a system name of your choice, or just hit enter to select the default.

#### 4.3.15 – tzsetup

Setup Time Zone

Time Zone (Argentina, Australia\_ACT, Australia\_Broken-Hill, Australia\_LHI, Australia\_NSW, Australia\_North, Australia\_Queensland, Australia\_South, Australia\_Sturt, Australia\_Tasmania, Australia\_Victoria, Australia\_West, Australia\_Yancowinna, Brazil\_Acre, Brazil\_DeNoronha, Brazil\_East, Brazil\_West, CET, Canada\_Atlantic, Canada\_Central, Canada\_East-Saskatchewan, Canada\_Eastern, Canada\_Mountain, Canada\_Newfoundland, Canada\_Pacific, Canada\_Yukon, Chile\_Continental, Chile\_EasterIsland, Cuba, EET, Egypt, GB-Eire, GMT, HST, Hongkong, Iceland, Iran, Israel, Jamaica, Japan, Libya, Mexico\_BajaNorte, Mexico\_BajaSur, Mexico\_General, NZ, NZ\_CHAT, Navajo, PRC, Poland, ROC, ROK, Singapore, Turkey, US\_Alaska, US\_Arizona, US\_Central, US\_East-Indiana, US\_Eastern, US\_Hawaii, US\_Michigan, US\_Mountain, US\_Pacific, US\_Yukon, W-SU, WET)[US\_Eastern]:

Daylight saving time (DST) is handled automatically. Type your chosen time zone and hit enter.

#### 4.3.16 – bootsetup

Setup Plan 9 FAT boot partition (9fat)

Plan 9 FAT partition (/dev/sdC0/9fat)[/dev/sdC0/9fat]:

Hit enter to accept the default. Any environment variables entered at the > prompt during boot, as well as settings configured during install will now be written to /n/9fat/plan9.ini and the kernel will be copied to the 9fat partition.

```
dosdrv: serving #s/dos
Initializing Plan 9 FAT partition.
% disk/format -r 2 -d -b /386/pbs /dev/sdC0/9fat
Initializing FAT file system
type hard, 12 tracks, 255 heads, 63 sectors/track, 512 bytes/sec
used 4096 bytes
% mount -c /srv/dos /n/9fat /dev/sdC0/9fat
% rm -f /n/9fat/9bootfat /n/9fat/plan9.ini /n/9fat/9pcf
% cp /n/newfs/386/9bootfat /n/9fat/9bootfat
% chmod +al /n/9fat/9bootfat
% cp /tmp/plan9.ini /n/9fat/plan9.ini
% cp /n/newfs/386/9pcf /n/9fat/9pcf
```

If you use the Windows NT/2000/XP master boot record or a master boot record from a Unix clone (e.g., LILO or FreeBSD bootmgr), it is probably safe to continue using that boot record rather than install the Plan 9 boot record.

Since we are not installing on a disk with a pre-existing Windows installation, we choose to install the Plan 9 master boot record and mark the partition active:

```
Install the Plan 9 master boot record (y, n)[no default]: y
Mark the Plan 9 partition active (y, n)[no default]: y
```

The Plan 9 partition is now marked as active.

#### 4.3.17 – finish

The final task is to remove (or disable) the CD-ROM and finish by hitting enter to reboot the machine.

Congratulations, you've just installed a 9front system!

Now your cat goes to sleep on the keyboard.



#### 4.4 – Encrypted Partitions



9front supports booting from encrypted `fsworm`, `fscache` and other partitions. The following assumes a hard drive `/dev/sdC0`:

- Begin the installation process as normal by booting `9front.iso`.
- During `partdisk` and `prepdisk`, create the `fsworm`, `fscache` and other partitions as normal.
- At the `mountfs` prompt, enter `!rc` to drop to a shell prompt.
- Format and activate the encrypted partitions (generates new encryption key):

```
cryptsetup -f /dev/sdC0/fsworm /dev/sdC0/fscache /dev/sdC0/other
cryptsetup -i /dev/sdC0/fsworm /dev/sdC0/fscache /dev/sdC0/other
```

- Type `exit` to resume the installation.
- During `mountfs`, select the partitions under `/dev/fs` instead of the ones under `/dev/sdC0`.
- During the rest of the installation, select the defaults, where appropriate.
- During post-installation boot, at the `bootargs` prompt, type `!rc` to drop to a shell prompt.

- Activate the encrypted partitions (prompts for password):

```
cryptsetup -i /dev/sdC0/fsworm /dev/sdC0/fscache /dev/sdC0/other
```

- Type `exit` to resume booting.
- At the `bootargs` prompt, enter `local!/dev/fs/fscache` to continue booting from the encrypted partition.

**Note:** Encrypted partitions are largely untested. Data integrity is not guaranteed. Use at your own risk.

Encrypted partitions have not been tested with `hjfs`.

Read: `cryptsetup(8)`

## FQA 5 – Building the System from Source



### 5.1 – Why should I build my system from source?

Some reasons why you might actually wish or need to build from source:

- Updates have been committed since you performed the installation.
- Test or develop new features.

### 5.2 – Building 9front from source

#### 5.2.1 – Update sources

9front uses hg(1) to synchronize the system with the 9front repository:

```
cd /
bind -ac /dist/plan9front /
hg incoming
hg -v pull -u
```

which is consolidated in the command:

```
sysupdate
```

If hg refuses to update due to conflicts (e.g. if you installed something locally which was later incorporated into 9front), this is a quick workaround which discards local changes:

```
hg update -C --clean
```

After the tree is updated, recompile/build the updated programs as needed.

### 5.2.1.1 – hgrc

To update when logged in as a user other than `glenda`, add the following to `$home/lib/hgrc`:

```
[trusted]
users=glenda
groups=glenda
```

To use Mercurial with `webfs(4)` instead of Python's built-in HTTP:

```
[extensions]
hgwebfs=
```

**Note:** The `hgwebfs` extension expects login credentials to be present in `factotum(4)`, else Mercurial will abort when any attempt is made to access resources that require authentication. To add credentials to `factotum`:

```
echo 'key proto=pass realm=PAIN server=code.9front.org service=http n
      user=username !password=FckG00gl!' >>/mnt/factotum/ctl
```

**Note:** Depending on who you are and the repository being used, your credentials may differ.

Credentials may be saved permanently and loaded into `factotum` at boot time using [FQA 8.4.7 – secstore](#).

### 5.2.2 – Building from source

**Note:** A minimum of 512MB RAM is needed to link some programs. If less than 512MB is available, be sure to turn on swap before building. Read: `swap(8)`, [FQA 3.3 – What is an appropriate first system to learn 9front on?](#)

```
# create any missing directories needed for the build
cd /
. /sys/lib/rootstub

# build everything
cd /sys/src
mk install
mk clean

# build manpage indices
cd /sys/man
mk

# build the papers and html (optional)
cd /sys/doc
mk
mk html
```

Build the kernel for 386:

```
cd /sys/src/9/pc
mk install
```

Build the kernel for amd64:

```
cd /sys/src/9/pc64
mk install
```

Build the kernel for Raspberry Pi:

```
cd /sys/src/9/bcm
mk install
```

Read: *FQA 7.2.5 – How do I install a new kernel?*

### 5.2.2.1 – Cross compiling

To cross compile, simply set the `objtype` environment variable prior to running the build. For example, to build all the amd64 binaries on a 386 system:

```
# create any missing directories needed for the build
cd /
. /sys/lib/rootstub
cd /sys/src
objtype=amd64 mk install
```

## 5.3 – Building an ISO

The 9front ISO is a livecd that also serves as install media.

**Note:** Currently, only the 386 architecture is built for the ISO. Read: *FQA 8.9 – Bootstrapping architectures not included on the ISO* for more information on booting other architectures.

```
# put your root file system into /n/src9
bind /root /n/src9

# put your hg repository there
bind -ac /dist/plan9front /n/src9

# build the iso
cd /sys/lib/dist
mk /tmp/9front.iso
```

## 5.4 – Common Problems when Compiling and Building

Most of the time, problems in the build process are caused by not following the above directions carefully.

## FQA 6 – Networking



### 6.1 – Before we go any further

Plan 9's approach to networking is unusual. For the bulk of this document, it helps if you have read and at least partially understood *FQA 0.1 – What is Plan 9?*

Read: *The Organization of Networks in Plan 9*

If you are working with applications such as web servers, FTP servers, and mail servers, you may benefit greatly by reading the RFCs. Most likely, you can't read all of them. Pick some topics that you are interested in, or that you use in your network environment. Look them up, find out how they are intended to work. The RFCs define many (thousands of) standards for protocols on the Internet and how they are supposed to work.

**Note:** A script for downloading all the RFCs is located in `/lib/rfc/grabrfc`. It may take hours for the script to run to completion.

### 6.2 – Network configuration

Normally, networking is initially configured by the installation process. However, it is good to understand what is happening in this process and how it works. In Plan 9, network configuration is organized in `ndb`, the network database.

From `ndb(6)` :

The network database consists of files describing machines known to the local installation and machines known publicly. The files comprise multi-line tuples made up of attribute/value pairs of the form `attr=value` or sometimes just `attr`. Each line starting without white space starts a new tuple. Lines starting with `#` are comments.



The file `/lib/ndb/local` is the root of the database. Other files are included in the database if a tuple with an attribute-value pair of attribute database and no value exists in `/lib/ndb/local`. Within the database tuple, each pair with attribute file identifies a file to be included in the database. The files are searched in the order they appear. For example:

```
database=  
  file=/lib/ndb/common  
  file=/lib/ndb/local  
  file=/lib/ndb/global
```

declares the database to be composed of the three files `/lib/ndb/common`, `/lib/ndb/local`, and `/lib/ndb/global`. By default, `/lib/ndb/local` is searched before the others. However, `/lib/ndb/local` may be included in the database to redefine its ordering.

Within tuples, pairs on the same line bind tighter than pairs on different lines.

### 6.2.1 – Host name

Each machine on the network receives a corresponding section in `ndb`. The host name (hereafter referred to as `sysname`) is assigned by setting the `sys=` tuple:

```
sys=x301
```

The resulting `sysname` is used by the `/rc/bin/termrc` and `/rc/bin/cpurc` startup scripts, which in turn call upon any additional configuration that may exist in `/cfg/$sysname/`. (Look at the scripts to see how they deal with `/cfg`.)

### 6.2.2 – Identifying and setting up your network interfaces

Network interfaces are recognized by their MAC addresses, which are identified to `ndb` using the `ether=` tuple:

```
sys=x301 ether=00226811f7dd
```

Additional tuples in the same grouping will be used to configure the interface in question.

#### 6.2.2.1 – Wifi

The following sections provide information pertaining to specific chipsets.

Read: `plan9.ini(8)`, *FQA Section 3.2 – Known Working Hardware*

##### 6.2.2.1.1 – Interfaces

### 6.2.2.1.1.1 – wavelan

Lucent Wavelan (Orinoco) IEEE 802.11b and compatible PCMCIA cards. Compatible cards include the Dell TrueMobile 1150 and the Linksys Instant Wireless Network PC Card. Port and IRQ defaults are 0x180 and 3 respectively.

These cards take a number of unique options to aid in identifying the card correctly on the 802.11b network. The network may be ad hoc or managed (i.e. use an access point): `mode=[adhoc, managed]` and defaults to managed. The 802.11b network to attach to (managed mode) or identify as (ad hoc mode), is specified by `ssid=string` and defaults to a null string. The card station name is given by `station=string` and defaults to Plan 9 STA. The channel to use is given by `channel=number` where number lies in the range 1 to 16 inclusive; the channel is normally negotiated automatically.

If the card is capable of encryption, the following options may be used: `crypt=[off, on]` and defaults to on. `keyN=string` sets the encryption key N (where N is in the range 1 to 4 inclusive) to string; this will also set the transmit key to N (see below). There are two formats for string which depend on the length of the string. If it is exactly 5 or 13 characters long it is assumed to be an alphanumeric key; if it is exactly 10 or 26 characters long the key is assumed to be in hex format (without a leading 0x). The lengths are checked, as is the format of a hex key. `txkey=number` sets the transmit key to use to be number in the range 1 to 4 inclusive. If it is desired to exclude or include unencrypted packets `clear=[off, on]` configures reception and defaults to inclusion.

The defaults are intended to match the common case of a managed network with encryption and a typical entry would only require, for example `ssid=left-  
armpit key1=afish key2=calledraawaru` if the port and IRQ defaults are used. These options may be set after boot by writing to the device's `ctl` file using a space as the separator between option and value, e.g. `echo 'key2  
1d8f65c9a52d83c8e4b43f94af' >/net/ether0/0/ctl` Card-specific power management may be enabled/disabled by `pm=[on, off]`

### 6.2.2.1.1.2 – wavelanpci

PCI Ethernet adapters that use the same Wavelan programming interface. Currently the only tested cards are those based on the Intersil Prism 2.5 chipset.

### 6.2.2.1.1.3 – iw1

Intel Wireless WiFi Link mini PCI-Express adapters require firmware from [http://firmware.openbsd.org/firmware/iwn-firmware\\*.tgz](http://firmware.openbsd.org/firmware/iwn-firmware*.tgz) to be present on attach in `/lib/firmware` or `/boot`. To select the access point, the `ssid=` and `bssid=` parameters can be specified at boot or set during runtime like:

```
echo ssid left-armpit >/net/ether1/clone
```

If both `ssid=` and `bssid=` are specified, both must match. Scan results appear in the `ifstats` file and can be read out like:

```
cat /net/ether1/ifstats
```

Ad-hoc mode or WEP encryption is currently not supported. To enable WPA/WPA2 encryption, see `wpa(8)` for details.

#### 6.2.2.1.1.4 – `rt2860`

Ralink Technology PCI/PCI-Express wireless adapters require firmware from [http://firmware.openbsd.org/firmware/ral-firmware\\*.tgz](http://firmware.openbsd.org/firmware/ral-firmware*.tgz) to be present on attach in `/lib/firmware` or `/boot`. See the `iwl` section above for configuration details.

#### 6.2.2.1.1.5 – `wpi`

Intel PRO Wireless 3945abg PCI/PCI-Express wireless adapters require firmware from [http://firmware.openbsd.org/firmware/\\*/wpi-firmware\\*.tgz](http://firmware.openbsd.org/firmware/*/wpi-firmware*.tgz) to be present on attach in `/lib/firmware` or `/boot`. See the `iwl` section above for configuration details.

#### 6.2.2.1.2 – `WPA`

WPA1/TKIP and WPA2/CCMP are supported with the use of the `wpa(8)` command.

Read: `wpa(8)`

#### 6.2.2.1.3 – `Wifi Roaming`

A script can be used to dynamically re-associate with available wifi access points:

<http://plan9.stanleylieber.com/rc/wifiroam>

Example usage:

```
@{wifiroam attwifi | aux/statusmsg -k wifiroam} &
```

#### 6.2.2.1.4 – `Wifi Debug`

For cards that use the `wifi` layer, debug prints (**note:** will appear on the console) may be enabled with:

```
echo debug >'#l0/ether0/clone'  
# change this to suit if wifi interface is not #l0
```

or by adding `debug=1` to the interface definition in `plan9.ini`.

Read: `plan9.ini(8)`

**6.2.3 – IP address** The `ip=` tuple is used to associate an IP address with the machine:

```
sys=x301 ether=00226811f7dd ip=192.168.0.31
```

If no `ip=` tuple is present, the boot scripts will attempt to bring up the interface using DHCP (see below).

## 6.2.4 – Default gateway

The default gateway is configured using the `ipgw=` tuple, usually under an `ipnet=` section that defines default settings for an entire subnet:

```
ipnet=9front ip=192.168.0.0 ipmask=255.255.255.0 ipgw=192.168.0.1
```

but it may also be specified on a per-machine basis:

```
sys=x301 ether=00226811f7dd ip=192.168.0.31 ipgw=192.168.0.1
```

**Note:** Tuples included in the definition of a machine supercede those defined for the network to which the machine belongs.

## 6.2.5 – DNS Resolution

DNS resolvers may be specified using the `dns=` tuple, and may be configured for an entire network:

```
ipnet=9front ip=192.168.0.0 ipmask=255.255.255.0 ipgw=192.168.0.1
      dns=192.168.0.1
```

or on a per-machine basis:

```
sys=x301 ether=00226811f7dd ip=192.168.0.31 dns=192.168.0.1
```

These changes will take effect after a reboot. To configure a DNS resolver on the fly, it is possible to manually edit `/net/ndb`:

```
ip=192.168.0.31 ipmask=255.255.255.0 ipgw=192.168.0.1
  sys=x301
  dom=x301.9front
  dns=192.168.0.1
  # add or modify dns= lines to associate the DNS
  # server 192.168.0.1 with the running system
```

**Note:** `/net/ndb` is a synthetic file that represents the current operating state. It does not persist across reboots.

Read: `ip(3)`

Finally, to turn on debug in `dns`:

```
echo -n debug >/net/dns
```

### 6.2.5.1 – Caching DNS server

To run a caching DNS server, add entries corresponding to the DNS rootservers in `/lib/ndb/local`:

```

#
# entries defining the dns root.  these will be overridden by any
# authentic info obtained from the root.
#
dom=
    ns=A.ROOT-SERVERS.NET
    ns=B.ROOT-SERVERS.NET
    ns=C.ROOT-SERVERS.NET
    ns=D.ROOT-SERVERS.NET
    ns=E.ROOT-SERVERS.NET
    ns=F.ROOT-SERVERS.NET
    ns=G.ROOT-SERVERS.NET
    ns=H.ROOT-SERVERS.NET
    ns=I.ROOT-SERVERS.NET
    ns=J.ROOT-SERVERS.NET
    ns=K.ROOT-SERVERS.NET
    ns=L.ROOT-SERVERS.NET
    ns=M.ROOT-SERVERS.NET

```

Next, modify `/cfg/$sysname/termrc` or `/cfg/$sysname/cpurc` (whichever is appropriate) to include the following:

```
ndb/dns -rs
```

The caching DNS server will be started at boot time.

Finally, modify `/lib/ndb/local` such that the desired machines will use the IP address of the new caching DNS server as their DNS server, either by changing the `dns=` tuple under the `ipnet` of the corresponding network or by adding a `dns=` tuple to the line of each desired machine.

Read: `ndb(6)` and `ndb(8)`

### 6.2.5.2 – DNS authoritative name server

An authoritative domain name record, with associated reverse-lookup and sub-domains, looks like this:

```

dom=bell-labs.co soa=
    refresh=10800 ttl=10800
    serial=2012110732          # automatically maintained if omitted
    ns=ns5.he.net
    ns=ns4.he.net
    ns=ns3.he.net
    ns=ns2.he.net
    ns=nm.iawtp.com
    ns=pp.iawtp.com
    ns=mars2.iawtp.com
    dnsslave=slave.dns.he.net
    mb=sl@stanleylieber.com
    mx=pp.inri.net pref=5
    mx=nm.inri.net pref=10
    mx=mars2.inri.net pref=15
    txtrr="v=spf1 mx -all"

dom=125.191.107.in-addr.arpa soa=
    refresh=3600 ttl=3600
    ns=nm.iawtp.com

dom=bell-labs.co ip=107.191.125.208

dom=www.bell-labs.co cname=bell-labs.co

```

An FQDN may be assigned to an existing machine by adding the `dom=` tuple to its definition:

```
sys=x301 dom=x301.bell-labs.co ether=00226811f7dd ip=192.168.0.31
```

**Note:** The `dnsslave` entries specify slave DNS servers that should be notified when the domain changes. The notification also requires the `-n` flag:

```
ndb/dns -nrs
```

Read: `ndb(8)`

## 6.2.6 – Network-wide configuration

Settings for an entire network subnet may be defined under an `ipnet=` tuple:

```

ipnet=9front ip=192.168.0.0 ipmask=255.255.255.0
    ipgw=192.168.0.1
    auth=192.168.0.2
    authdom=9front
    fs=192.168.0.3
    cpu=192.168.0.4
    dns=192.168.0.1
    dnsdom=9front
    smtp=192.168.0.4

# ethernet/wifi router
sys=onoff dom=onoff.9front ip=192.168.0.1

# auth server
sys=auth dom=auth.9front ether=00d059b6dac8 ip=192.168.0.2
    bootf=/386/9bootpxe

# cpu server
sys=cpu dom=cpu.9front ether=001125149137 ip=192.168.0.4
    bootf=/386/9bootpxe

# file server
sys=fs dom=fs.9front ether=001641360117 ip=192.168.0.3

# terminal
sys=x301 dom=x301.9front ether=00226811f7dd ip=192.168.0.31
    bootf=/386/9bootpxe

```

## 6.2.7 – Activating the changes

### 6.2.7.1 – NIC

Network interfaces are automatically initialized at boot time. To make a manual change without rebooting, use the `ipconfig(8)` command:

```
ip/ipconfig -g 192.168.0.1 ether /net/ether0 192.168.0.31 255.255.255.0
```

### 6.2.7.2 – cs

To refresh the network database **NOW** after changing `/lib/ndb/local`:

```
echo -n refresh > /net/cs
```

### 6.2.7.3 – dns

```
echo -n refresh > /net/dns
```

## 6.2.8 – Verifying network settings

```
% cat /net/ndb
ip=192.168.0.31 ipmask=255.255.255.0 ipgw=192.168.0.1
  sys=x301
  dom=x301.9front
  auth=192.168.0.2
  dns=192.168.0.1
```

### 6.2.8.1 – Checking routes

```
% cat /net/iproute
0.0.0.0          /96  192.168.0.1      4    none  -
192.168.0.0     /120 192.168.0.0      4i   ifc   0
192.168.0.0     /128 192.168.0.0      4b   ifc   -
192.168.0.31    /128 192.168.0.31     4u   ifc   0
192.168.0.255   /128 192.168.0.255    4b   ifc   -
255.255.255.255 /128 255.255.255.255  4b   ifc   0
```

#### 6.2.8.1.1 – Adding static routes

Route requests for 192.168.1.0/24 through the gateway 192.168.0.99 (which itself must already be accessible via the existing network configuration):

```
echo 'add 192.168.1.0 255.255.255.0 192.168.0.99' >/net/iproute
```

**Note:** Manual configuration such as this may be added to optional boot scripts created in /cfg/\$sysname/.

Read: ip(3)

## 6.2.9 – Setting up your 9front box as a forwarding gateway

Read: ip(3)

### 6.2.10 – Setting up aliases on an interface

Read: ip(3)

## 6.3 – How do I filter and firewall with 9front?

No.

## 6.4 – Dynamic Host Configuration Protocol (DHCP)



### 6.4.1 – DHCP client

In `/lib/ndb/local`, if no `ip=` tuple is present in the machine's definition, the boot scripts will attempt to obtain an IP address via DHCP.

To obtain a DHCP lease manually:

```
ip/ipconfig
```

Read: `ipconfig(8)`

### 6.4.2 – DHCP server

From `dhcpcd(8)` :

Dhcpd runs the BOOTP and DHCP protocols. Clients use these protocols to obtain configuration information. This information comes from attribute/value pairs in the network database (see ndb(6) and ndb(8)). DHCP requests are honored both for static addresses found in the NDB and for dynamic addresses listed in the command line. DHCP requests are honored if either:

- there exists an NDB entry containing both the ethernet address of the requester and an IP address on the originating network or subnetwork.
- a free dynamic address exists on the originating network or subnetwork.

A BOOTP request is honored if all of the following are true:

- there exists an NDB entry containing both the ethernet address of the requester and an IP address on the originating network or subnetwork.
- the entry contains a bootf= attribute
- the file in the bootf= attribute is readable.

Dynamic addresses are specified on the command line as a list of addresses and number pairs. For example,

```
ip/dhcpd 10.1.1.12 10 10.2.1.70 12
```

directs dhcpd to return dynamic addresses 10.1.1.12 through 10.1.1.21 inclusive and 10.2.1.70 through 10.2.1.81 inclusive.

Dhcpd maintains a record of all dynamic addresses in the directory /lib/ndb/dhcp, one file per address. If multiple servers have access to this common directory, they will correctly coordinate their actions.

Attributes come from either the NDB entry for the system, the entry for its subnet, or the entry for its network. The system entry has precedence, then the subnet, then the network. The NDB attributes used are:

ip	the IP address
ipmask	the IP mask
ipgw	the default IP gateway
dom	the domain name of the system
fs	the default Plan 9 file server
auth	the default Plan 9 authentication server
dns	a domain name server
ntp	a network time protocol server
time	a time server
wins	a NETBIOS name server
www	a World Wide Web proxy
pop3	a POP3 mail server
smtp	an SMTP mail server
bootf	the default boot file; see ndb(6)

Dhcpd will answer BOOTP requests only if it has been specifically targeted or if it has read access to the boot file for the requester. That means that the requester must specify a boot file in the request or one has to exist in NDB for dhcpd to answer. Dhcpd will answer all DHCP requests for which it can associate an IP address with the requester.

To configure a DHCP server on your system, modify `/cfg/$sysname/cpurc` or `/cfg/$sysname/termrc` (whichever is appropriate) to start `dhcpcd` and `tftpd` at boot time:

```
ip/dhcpcd
```

Read: `dhcpcd(8)`

## 6.5 – PPP

Read: Dailup modem config at the Bell Labs Plan 9 wiki.

## 6.6 – Setting up a network bridge in 9front

Read: `bridge(3)` and `ip(3)`

## 6.7 – How do I boot from the network?

First, read *FQA 7.3.3 – Setting up a listener for network connections*. The file server should already be running a listener, and an auth server should already be configured and running on the network.

### 6.7.1 – How do I tcp boot?

It is possible to boot from local media and then mount the root file system over the network. At the `bootargs` prompt, type `tcp`. At this point, `ip/ipconfig` will determine network parameters using DHCP. When file (`fs`) or authentication (`auth`) server IP addresses could not be determined over DHCP then the boot process will prompt for those. When prompted for a `user`, enter a valid username and password that has already been configured on the auth server. The machine should then proceed to mount its root file system from the file server.

**Note:** Values for `fs` and `auth` may be added to `plan9.ini`.

#### 6.7.1.1 – Passing arguments to ipconfig at the bootargs prompt

When a DHCP server is not available, you may still `tcp` boot by configuring networking manually at the `bootargs` prompt. Everything after `tcp!` is passed as arguments to the `ipconfig` command.

At the prompt:

```
bootargs is (tcp, il, local!device) [tcp]
```

enter something like the following:

```
tcp!-g 192.168.0.1 ether /net/ether0 192.168.0.23 255.255.255.0
```

where `192.168.0.1` is the gateway, `192.168.0.23` is the static IP address and `255.255.255.0` the subnet mask.

Read: `ipconfig(8)`

## 6.7.2 – How do I boot using PXE?

It is also possible to PXE boot a system.

On the file server, add the following lines to `/cfg/$sysname/cpurc` to start `dhcpd` and `tftpd` at boot time:

```
ip/dhcpd
ip/tftpd
```

Add an entry for `tftp` under the appropriate `ipnet=` tuple in `/lib/ndb/local`:

```
ipnet=9front ip=192.168.0.0 ipmask=255.255.255.0 ipgw=192.168.0.1
  auth=192.168.0.2
  authdom=9front
  cpu=192.168.0.4
  dns=192.168.0.1
  dnsdom=9front
  smtp=192.168.0.4
  tftp=192.168.0.3
```

Reboot the file server:

```
fshalt -r
```

To configure machines that will PXE boot from the file server, edit `/lib/ndb/local` on the file server and add a `bootf=` (boot file) tuple to the line representing each machine:

```
sys=x301 dom=x301.9front ether=00226811f7dd ip=192.168.0.31
  bootf=/386/9bootpxe
```

For the system `x301` we would then create a file `/cfg/pxe/00226811f7dd` on the file server to serve as its `plan9.ini`:


```
bootfile=/386/9pcf
bootargs=tcp
nbootprompt=tcp
auth=192.168.0.2
fs=192.168.0.3
mouseport=ps2intellimouse
monitor=vesa
vgasize=1440x900x32
*acpi=1
user=sl
```

**Note:** The `user=` parameter refers to a single username that has been added both to the file server (for file permissions) as well as to the auth server (for network authentication).

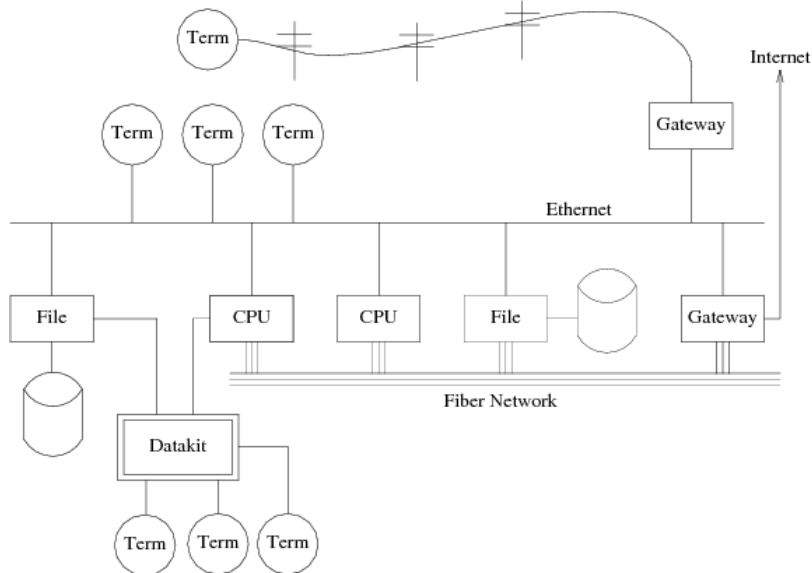
If a file matching the remote system's MAC address is not found under `/cfg/pxe/`, the file `/cfg/pxe/default` (if it exists) will be used instead.

Finally, boot the desired remote systems via PXE. When prompted for a user, enter a valid username and password that has already been configured on the auth server. The remote system should now proceed to boot from the file server.

## FQA 7 – System Management



### 7.1 – Plan 9 Services Overview



In order to be an effective system manager it is a good idea to understand how the system is designed, and how it is intended to be used.

Plan 9 services may be run on separate machines, all together on one machine, or in various combinations. The original design of Plan 9 assumed that each network service would run on separate hardware; by design, individual components of the system are generally unaware if they co-exist on the same machine or are distributed amongst

separate machines.

Read: *Designing Plan 9, Plan 9 From Bell Labs, The Organization of Networks in Plan 9*

### 7.1.1 – What is the kernel?

The kernel is a service that provides processes and resources for users active on the current system. Every Plan 9 system boots a kernel.

At boot time the kernel takes on the identify of `$user` (the user who logs in at the console), which becomes the `hostowner` of the system. The `hostowner` in turn 1.) controls access to the kernel's resources, 2.) serves as the auth identity (`authid`) of the machine and the services it provides.

**Note:** The `hostowner` differs from the concept of `root` on a UNIX system, where a single user `root` may take control of all processes *and* files on the system. By contrast, even the `hostowner` of a Plan 9 file server cannot violate file permissions on the file server, except when permissions checking is disabled on the console or when entering special commands at the console of the file server. The `hostowner` controls only the *processes* running on the local machine. This fundamental separation between control of processes and file permissions is exploited throughout the system, but can be confusing for users coming from a UNIX background.

### 7.1.2 – What is the file server?

In a traditional Plan 9 network there is one file server, typically the only machine with a physical hard disk, that serves files to all other machines on the network. In most cases, other machines are either diskless or only use their disks for local caching. In modern Plan 9, the file server typically runs with `service=cpu` in `plan9.ini`, and runs the `/rc/bin/cpure` script at boot time. It stores its `authid/authdom/secstore/password` in a file `/dev/sdXX/nvram`, which corresponds with the values for that user (the `authid`) held by the auth server. Users seeking access to the file server must be added as a user on the file system itself, and, if auth is enabled, added to the auth server's user database.

9front supports two different file systems for use on the file server: `cwfs` and `hjfs`. `cwfs` is a userspace port of Ken Thompson's original Plan 9 file server, which originally ran a unique kernel and *only* served files. `hjfs` is a new, experimental file server that stores both the cache and worm on a single partition (and thus requires less disk space to be used effectively).

Read: *The Plan 9 File Server* (deprecated, but partially applies to `cwfs`), `cwfs(4)`, `hjfs(4)`

### 7.1.3 – What is the auth server?

The auth server manages authentication for an entire Plan 9 network. It runs with `service=cpu` in `plan9.ini`, and runs the `/rc/bin/cpure` script at boot time. When a disk is present, it stores its `authid/authdom/secstore/password` in a file `/dev/sdXX/nvram`. If no `nvram` is found, the `auth/wrkey` command is run automatically at boot time.

**Note:** The `cron(8)` service should be run only on the auth server, where it can authenticate itself to access any of the other machines on the network.

Read: *Security in Plan 9*, `auth(8)`

#### 7.1.4 – What is the cpu server?

The `cpu` server is used for remote computation. It runs with `service=cpu` in `plan9.ini`, and runs the `/rc/bin/cpurc` script at boot time. When a disk is present, it stores its `authid/authdom/secstore/password` in a file `/dev/sdXX/nvram`, which corresponds with the values for that user (the `authid`) held by the auth server. If no `nvram` is found, the `auth/wrkey` command is run automatically at boot time. Users seeking access to a `cpu` server must first be added on the file system of the `cpu` server's corresponding file server (for permission to access and modify files) as well as the user database of its auth server (for login authentication). In practice, the file server and auth server (and even some terminals) will often run `cpu` listeners to enable remote command line access.

The common use case for a `cpu` server is to execute programs compiled for a different architecture than that of the terminal, or to execute programs closer to the data they are operating on (for example, if the terminal is running over a slow link but the `cpu` server is on the same ethernet segment as the file server). In the early days of Plan 9, a `cpu` server was often significantly more powerful than the (often, special-purpose) hardware used for diskless terminals. Today, terminals are typically powerful computers in their own right, and the need for a separate machine running only as a `cpu` server is less common. That said, it can be useful to execute unstable or unpredictable programs on a separate machine so that rebooting does not affect one's immediate workspace environment — especially when testing new code.

Read: *The Organization of Networks in Plan 9*, `cpu(1)`, `exportfs(4)`

#### 7.1.5 – What is a terminal?

The terminal is the machine at which the Plan 9 user is most often physically located. Usually diskless, it runs with `service=terminal` in `plan9.ini` and executes the `/rc/bin/termrc` script at boot time. Terminal users who do not run stand alone machines or who wish to access Plan 9 network resources must first be added to the file system of the network's file server and to the user database of the network's auth server.

**Note:** Many Plan 9 users run stand alone systems that operate — effectively — as a combined terminal and file server. For example, inside a virtual machine such as `qemu`, or booted from hard disk on a laptop. In this case the Plan 9 network is entirely self-contained, running one kernel on one machine, which renders `auth` and `cpu` services superfluous. This configuration trades some of the inherent security of separate hardware and kernel boundaries for the convenience of combining the whole system into a single, bootable instance.

## 7.2 – Kernel configuration and maintenance

### 7.2.1 – How do I mount the 9fat partition?

9front has done away with the scripts `9fat:`, `c:`, and so forth, that are found in the Bell Labs Plan 9 distribution. Instead, use the `9fs` script to mount the 9fat partition:

```
9fs 9fat
```

If you are not at the console, or if `#S` has not already been bound over `/dev`:

```
bind -b '#S' /dev # bind the local hard drive kernel device over /dev
9fs 9fat /dev/sdXX/9fat # specify the full path to the corresponding 9fat
```

**Note:** `9fs 9fat` posts a file descriptor in `/srv/dos`. If this file already exists and is already in use, `9fs 9fat` will fail. If no other process is using the file it is safe to simply remove it and run `9fs 9fat` again.

Read: `dosrv(4)`

### 7.2.2 – How do I modify `plan9.ini`?

Mount the 9fat partition and then edit the file `/n/9fat/plan9.ini`.

**Note:** The file must end with a newline.

Read: `plan9.ini(8)`

### 7.2.3 – Kernel configuration file

Kernel configuration files are stored in the kernel directory and share the name of the kernel to which they apply. For example, the configuration file for the `pcf` kernel is `/sys/src/9/pc/pcf`.

### 7.2.4 – Kernel drivers

Kernel driver source files are located in the kernel source directory.

### 7.2.5 – How do I install a new kernel?

First, build the new kernel(s).

For 386:

```
cd /sys/src/9/pc
mk install # kernel
```

For amd64:



```
cd /sys/src/9/pc64
mk install # kernel
```

Next, install the new bootloader and kernels into the 9fat partition. **Note:** The bootloader needs to be continuous on disk, so simply copying over the original file does not produce the desired effect. Instead:

```
9fs 9fat
rm /n/9fat/9bootfat
cp /386/9bootfat /n/9fat/
chmod +a1 /n/9fat/9bootfat # defrag magic
```

then copy the desired kernels:

For 386:

```
cp /386/9pcf /n/9fat/
```

For amd64:

```
cp /amd64/9pc64 /n/9fat/
```

## 7.3 – Fileserver configuration and maintenance

### 7.3.1 – Adding users

Add a new user on the file server:

For cwfs:

```
echo newuser username >>/srv/cwfs.cmd
```

For hjfs:

```
echo newuser username >>/srv/hjfs.cmd
```

If needed, make the new user a member of the upas (email) group:

For cwfs:

```
echo newuser upas +username >>/srv/cwfs.cmd
```

For hjfs:

```
echo newuser upas +username >>/srv/hjfs.cmd
```

In all cases, the file server's user database is stored in `/adm/users`. Examine this file, and the contents of the `/usr` directory, to evaluate success.

**Note:** It is also possible to access the control file interactively by running the command:

```
con -C /srv/cwfs.cmd
```

From here commands may be entered directly.

Type `Ctrl-\` to resume the `con` prompt, followed by `q` to quit.

The `newuser` file system command is described in the man pages `fs(8)` (for `cwfs`) and `hjfs(8)`

**Note:** New users are created without a profile, mail directory, `tmp` directory (needed to edit files with `sam`) or other confections. To install a default profile for a new user, upon first login as that user, run:

```
/sys/lib/newuser
```

then edit `/usr/username/lib/profile` to your own specifications.

### 7.3.2 – Configuring `nvr`

The `cpu` kernel checks the `nvr` file for valid auth credentials and attempts to copy them into `factotum` so that the machine may boot without manual intervention. To configure the `nvr`, run the command `auth/wrkey`, which will prompt for an `authid`, `authdom`, `secstore key`, and `password`. The `authid` is a synonym for the `hostowner` of the machine and should be a valid user that has already been (or will be) added to the corresponding auth server, in this case `glenda`. The `authdom` is the authentication domain for the machine, in this case `9front`. The `secstore key` and `password` are secret passwords of eight characters or more in length. The `password` is the password belonging to the `authid` user on the auth server responsible for the `authdom` entered above. The `secstore key` is the password of the user on the `secure-store` server. Read: `secstore(1)` and `secstore(8)`

Run the command `auth/wrkey`:

```
bad nvr key
bad authentication id
bad authentication domain
authid: glenda
authdom: 9front
secstore key: [glenda's password]
password: [glenda's password]
```

**Note:** Booting the file system with authentication enabled and an invalid `nvr` file will cause `auth/wrkey` to be run automatically at startup.

Read: `auth(8)`

### 7.3.3 – Setting up a listener for network connections

In order for remote machines to mount the file system of the file server it must first be running a network listener.

The first step is to switch from the terminal service to the `cpu` service by editing the service line in `/n/9fat/plan9.ini`:

```
service=cpu
```

Read: *FQA 7.2.2 – How do I modify plan9.ini?*

Before rebooting, configure the nvram: *FQA 7.3.2 – Configuring nvram*. This allows the machine to load auth credentials from the nvram file into factotum without manual intervention.

Reboot:

```
fshalt -r
```

The next step (on cwfs; not needed on hjfs) is to enable auth on the file server, to prevent unauthorized users from accessing the disk over the network. At the bootargs prompt, retype the default and add the `-c` flag to enter the file server's config mode. At the config prompt, type `noauth` twice to toggle authentication on the file server. Finally, type `end` to continue with the boot process:

```
bootargs is (tcp, local!device) [local!/dev/sdXX/fscache] local!/dev/sdXX/fscach
config: noauth
auth is now disabled
config: noauth
auth is now enabled
config: end
```

The machine will now continue to boot.

Once booted, the next step is to configure the file server to listen for connections from remote hosts. Modify the bootargs of the file server in `/n/9fat/plan9.ini`:

For cwfs:

```
bootargs=local!/dev/sdXX/fscache -a tcp!*!564
```

For hjfs:

```
bootargs=local!/dev/sdXX/fs -m 702 -A -a tcp!*!564
```

**Note:** The `-m 702` flag for `hjfs` allocates 702 megabytes of memory to be used as a cache. This value is typically automatically calculated by the 9front installer, and may differ on your system. There is no need to change whatever default was already configured.

Reboot the file server:

```
fshalt -r
```

When the system finishes booting it should now be listening for network connections to the file system. Users who have been added to the file server and the auth server should now be able to authenticate and mount the file server (tcp boot, etc.).

Read: `cwfs(4)`, `hjfs(4)`, *FQA 6.7.1 – How do I tcp boot?*

## 7.4 – Auth server configuration and maintenance

### 7.4.1 – Configuring an auth server

The auth server should be booted with `service=cpu` in `plan9.ini`, and `ndb` modified to associate the new auth server with the desired `authdom`.

If the `cpu` server machine boots from a local disk, edit the `service` line in in `/n/9fat/plan9.ini`:

```
service=cpu
```

Read: *FQA 7.2.2 – How do I modify plan9.ini?*

If the machine boots via PXE, edit the `service` line in in the file under `/cfg/pxe/` that correspondes to its MAC address. In this case, `/cfg/pxe/000c292fd30c`:

```
service=cpu
```

**Note:** The contents of `/cfg/pxe/000c292fd30c` serves as the equivalent of `plan9.ini` for the PXE booted machine. Any other settings that would normally be configured in `plan9.ini` may also be entered there.

Next, `ndb` must be modified to associate the new auth server with the desired `authdom`. Assuming the auth server has a MAC address of `00:0c:29:2f:d3:0c`, an IP address of `192.168.0.2`, and a default gateway/DNS server of `192.168.0.1` that are all on the Class C network `192.168.0.0/24`, and that the `authdom` is `9front`, edit `/lib/ndb/local` and add the `authdom` and the auth server's IP under the corresponding `ipnet`:

```
ipnet=9front ip=192.168.0.0 ipmask=255.255.255.0
  ipgw=192.168.0.1
  auth=192.168.0.2 # add auth server's ip
  authdom=9front # add authdom
  fs=192.168.0.3
  cpu=192.168.0.4
  dns=192.168.0.1
  dnsdom=9front
  smtp=192.168.0.4
```

Read: `ndb(6)`

Before rebooting, configure the `nvr`am: *FQA 7.3.2 – Configuring nvr*am. This allows the machine to load auth credentials from the `nvr`am file into `factotum` without manual intervention.

**Note:** If the auth server's `hostowner` (referred to as `authid` in the `auth/wrkey` dialogue) will be any other user than the default `glenda`, that user must be authorized (in the `auth` context) to "speak for" other users. Assuming a `hostowner` of `s1`, add a rule to `/lib/ndb/auth`:

```
hostid=s1
  uid=!sys uid=!adm uid=*
```

Read: `auth(8)`

Reboot:

```
fshalt -r
```

At boot time, the shell script `/rc/bin/cpure` consults `ndb` to determine if the machine is an auth server. If it is, the script will launch the `keyfs` process and start listeners for auth connections. If, after booting, `keyfs` is not running, something went wrong.

Finally, create an auth user and configure an auth password for the hostowner of the machine. This auth user should be the same name as the `authid` that was entered at boot time during the `auth/wrkey` dialogue. Likewise, set the `password` to match the `password` that was entered during the `auth/wrkey` dialogue.

Read: *FQA 7.4.2 – Adding users*

#### 7.4.1.1 – Avoiding an `ndb` entry for the auth server

If an auth server for a given `authdom` is not found in the local `ndb`, then the `authdial()` function from the `libauthsrv` library (used for resolving auth servers) will default to the dns host name `p9auth.example.com`, where `p9auth` is the sub-domain, and `example.com` is the `authdom`. This convention (where followed) is useful to avoid having to manually add auth server information for arbitrary remote networks to the local `ndb`.

#### 7.4.2 – Adding users

To add a new user to the auth server, login as the auth server's `hostowner`, make sure `auth/keyfs` is running in your namespace, and then set an auth password for the user:

```
auth/changeuser username
Password: # type password here, will not echo
Confirm password: # confirm password here, will not echo
assign Inferno/POP secret? (y/n) n
Expiration date (YYYYMMDD or never)[return = never]:
2 keys read
Post id:
User's full name:
Department #:
User's email address:
Sponsor's email address:
user username installed for Plan 9
```

**Note:** Questions that appear after the `keys read` notice are optional. Hit `Enter` for each one to leave them blank.

Read: `auth(8)`, `keyfs(4)`

### 7.4.3 – secstored

Secstore authenticates to a secure-store server using a password and optionally a hardware token, then saves or retrieves a file. This is intended to be a credentials store (public/private keypairs, passwords, and other secrets) for a factotum.

To set up secstored, login to the auth server as hostowner and:

```
mkdir /adm/secstore
chmod 770 /adm/secstore
```

Start secstored at boot time by adding the following to /cfg/\$sysname/cpurc on the auth server:

```
auth/secstored
```

Read: secstore(1), secstore(8)

#### 7.4.3.1 – Adding users to secstore

secuser is an administrative command that runs on the secstore machine, normally the auth server, to create new accounts and to change status on existing accounts. It prompts for account information such as password and expiration date, writing to /adm/secstore/who/user for a given secstore user.

Login to the auth server as hostowner and:

```
auth/secuser username
```

and answer the prompts.

By default, secstored warns the client if no account exists. If you prefer to obscure this information, use secuser to create an account FICTITIOUS.

Read: *FQA 8.4.7 – secstore* for more information on actually *using* secstore.

## 7.5 – Cpu server configuration and maintenance

### 7.5.1 – Configuring a cpu server

**Note:** Operating a cpu server requires auth services. Read: *FQA 7.4 – Auth server configuration and maintenance*

The first step in converting a terminal to a cpu server is to switch from the terminal service to the cpu service.

If the cpu server machine boots from a local disk, edit the service line in in /n/9fat/plan9.ini:

```
service=cpu
```

Read: *FQA 7.2.2 – How do I modify plan9.ini?*

If the machine boots via PXE, edit the `service` line in the file under `/cfg/pxe/` that corresponds to its MAC address. In this case, `/cfg/pxe/000c292fd30c`:

```
service=cpu
```

**Note:** The contents of `/cfg/pxe/000c292fd30c` serves as the equivalent of `plan9.ini` for the PXE booted machine. Any other settings that would normally be configured in `plan9.ini` may also be entered here.

Setting `service=cpu` causes the shell script `/rc/bin/cpurc` to be run at boot time, which in turn launches a listener that scans the `/rc/bin/service` directory for scripts corresponding to various network ports. Read: `listen(8)`. The script `tcp17010` handles incoming `cpu` connections. Authentication for incoming `cpu` connections is performed by the auth server associated with the `authdom` by `ndb`. Read: *FQA 7.4.1 – Configuring an auth server*

Before rebooting, configure the `nvr`am: *FQA 7.3.2 – Configuring nvr*am. This allows the machine to load auth credentials from the `nvr`am file into `factotum` without manual intervention.

Reboot:

```
fshalt -r
```

Read: *FQA 7.6.2 – Configuring a Terminal to Accept cpu Connections*

## 7.6 – Terminal configuration and maintenance

### 7.6.1 – Configuring a terminal

The `9front` ISO boots into a `livecd` running the `9pcf` kernel, resulting in the simplest form of terminal running on the 386 architecture. A terminal may also be network booted (the preferred method) or installed to its own stand-alone file system on a local storage device.

Read: *FQA 6.7 – How do I boot from the network?*

### 7.6.2 – Configuring a Terminal to Accept cpu Connections

If the `hostowner` `factotum` has been loaded with the appropriate key and the system is listening for `cpu` connections, a user may `cpu` into a terminal that is not running auth services. To configure a terminal to accept `cpu` connections in this fashion, substitute your choice of `dom` (this refers to the `authdom`), `user` and `password`, below:

```
echo 'key proto=p9sk1 dom=9front user=glenda !password=p@ssw0rd' \  
>/mnt/factotum/ctl  
aux/listen1 -t tcp!*!17010 /bin/cpu -R &
```

### 7.6.3 – UTC Timesync

By default, `/rc/bin/termrc` sets `TIMESYNCARGS=(-rLa1000000)`, to synchronize 9front time with the real time clock. On many systems this time is saved as UTC, whereas Windows keeps the local time there. If your time is in UTC you should omit the `-L`: Put `TIMESYNCARGS=(-ra1000000)` into `/rc/bin/termrc.local`, which gets included by `/rc/bin/termrc`.

## 7.7 – Mail server configuration and maintenance



Incoming and outgoing mail is handled by `upas` and its related suite of programs. Configuration is handled by a number of files located in `/mail/lib/`.

**Note:** Be sure to configure proper DNS entries for your domains. If Plan 9 will host your DNS, see: *FQA 6.2.5.2 – DNS authoritative name server*

Read: *Upas – A Simpler Approach to Network Mail*, `mail(1)`

### 7.7.1 – `smtpd.conf`

Some changes to the default `smtpd.conf` are required to accept mail for Internet domain names and to send mail for other hosts. The following lines should be changed to correspond to your network:



```
# outgoing mail will be sent from this domain by default
defaultdomain          9front.org

# do not be an open relay
norelay                on

# disable dns verification of sender domain
verifysenderdom       off

# save blocked messages
saveblockedmsg         off

# if norelay is on, you need to set the
# networks allowed to relay through
# as well as the domains to accept mail for
ournets 199.191.58.37/32 199.191.58.42/32 192.168.4.0/24

# domain names for which incoming mail is accepted
ourdomains 9front.org, bell-labs.co, cat-v.org
```

Read: smtpd(6), smtp(8)

### 7.7.2 – rewrite

To act as an Internet mail server, copy `rewrite.direct` to `rewrite` and modify accordingly:

```
# case conversion for postmaster
pOsTmAsTeR alias postmaster

# local mail
\!(.*) alias \1
(ttr|9front.org|bell-labs.co|cat-v.org)!(.*)    alias \2
[^\!@]+ translate "/bin/upas/aliasmail '&'"
local!(.*) >> /mail/box/1/mbox

# we can be just as complicated as BSD sendmail...
# convert source domain address to a chain a@b@c@d...
@[^\!@,]*):([^\!@]*)@([^\!@]*) alias \2@\3@\1
@[^\!@]*),([^\!@,]*):([^\!@]*)@([^\!@]*) alias @\1:\3@\4@\2

# convert a chain a@b@c@d... to ...d!c!b!a
([^\!@]+)@([^\!@]+)@(.+) alias \2!\1@\3
([^\!@]+)@([^\!@]+) alias \2!\1

# /mail/lib/remotemail will take care of gating to systems we don't know
([^\!@]*)!(.*) | "/mail/lib/qmail '\s' 'net!\1'" ""\2''
```

Read: `rewrite(6)`

### 7.7.3 – names.local

To map incoming e-mail addresses to local usernames, edit `names.local` accordingly:

```
# postmaster goes to glenda
postmaster          glenda
```

**Note:** *postmaster@[any domain]* will be delivered to local user *glenda*.

#### 7.7.4 – remotemail

Finally, *upas* needs to know what to do with mail that cannot be delivered locally. Edit *remotemail* and enter the desired behavior.

To deliver mail directly to the remote server responsible for the domain name in question:

```
#!/bin/rc
shift
sender=$1
shift
addr=$1
shift
exec /bin/upas/smtp $addr $sender $*
```

Read: *smtp(8)*

#### 7.7.5 – SMTP over TLS

First, make sure you have already created TLS certificates for your server.

Next, create a file */rc/bin/service/tcp587*:

```
#!/bin/rc
user='{cat /dev/user}'
exec /bin/upas/smtpd -c /sys/lib/tls/cert -n $3
# to use with listen1, change $3 to $net
```

#### 7.7.6 – IMAP4 over TLS

First, make sure you have already created TLS certificates for your server.

Next, create a file */rc/bin/service/tcp993*:

```
#!/bin/rc
exec tlssrv -c/sys/lib/tls/cert -limap4d \
  -r'{cat $3/remote} /bin/ip/imap4d -p \
  -r'{cat $3/remote} >>[2]/sys/log/imap4d
# to use with listen1, change $3 to $net
```

#### 7.7.7 – Spam Filtering

##### 7.7.7.1 – ratfs

From *ratfs(4)*:

*Ratfs* starts a process that mounts itself (see *bind(2)*) on mountpoint (default */mail/ratify*). *Ratfs* is a persistent representation of the local network configuration and spam blocking list. Without it each instance of *smtpd(6)* would need to reread and parse a multimegabyte list of addresses and accounts.

To configure the spam blocking list, edit `/mail/lib/blocked` as desired, according to the rules laid out in the man page. Example:

```
# allow messages from any user at 9front.org
*allow 9front.org!*

# block messages from any user at bell-labs.com
*block bell-labs.com!*

# block messages from ip block of aol modems
block 152.166.0.0/15
```

If `ratfs` is already running, cause it to reload the modified `/mail/lib/blocked`:

```
echo reload >/mail/ratify/ctl
```

For more details, read: `ratfs(4)`, `smtpd(6)`

To launch `ratfs` at boot time, add the following line to `/cfg/$sysname/cpustart`:

```
upas/ratfs
```

and add the following line to `/lib/namespace`:

```
mount -c #s/ratify /mail/ratify
```

**Note:** The `tcp25` and/or `tcp587` listener needs to be started *after* the modifications above so that `ratfs` will be accessible from the listener's namespace.

### 7.7.7.2 – scanmail

Read: `scanmail(8)`

## 7.8 – Web server configuration and maintenance

### 7.8.1 – ip/httpd

No.

### 7.8.2 – rc-httpd

The `rc-httpd` web server is a simple shell script that handles static files, directory listings and drop-in CGI programs such as the `werc` anti-framework. `rc-httpd` is run from a file in the directory scanned by `listen(8)`, or called as an argument to `listen1(8)`.

Read: `rc-httpd(8)`

## 7.9 – TLS certificates

To use TLS-enabled services on a Plan 9 mail server (`poptls`, `apoptls`, `imaps`, etc.) you need to generate a certificate and key for your mail server and tell the factotum of the server about that key:


```
ramfs -p
cd /tmp
auth/rsagen -t 'service=tls role=client owner=*' > key
chmod 600 key
cp key /sys/lib/tls/key # or: store key in secstore
auth/rsa2x509 'C=FR CN=*.fakedom.dom' /sys/lib/tls/key | \
    auth/pemencode CERTIFICATE > /sys/lib/tls/cert
```

where FR is a two digit country code, and fakedom.dom is the fully qualified domain-name.)

To load the key into the server's factotum at boot time, add the following line to /cfg/\$sysname/cpustart:

```
cat /sys/lib/tls/key >>/mnt/factotum/ctl
```

## FQA 8 – Using 9front



*i am form china  
here is lots of people knows  
and documented is chinese  
the time comllexity is need you test yourself  
and you will find it  
— wenwei peng*

When applied consistently, simple conventions can combine to provide useful results. In Plan 9, *conventions* are preferred to *rules*. This section explores the Plan 9 approach to actually using the computer.

### 8.1 – rc

The `rc` shell was written by Tom Duff for Research UNIX v10. It was later adopted as the shell for Plan 9. Some of its conventions are unusual compared with other command interpreters influenced by the Bourne shell. Although its syntax may seem strange at first, have patience; `rc` was designed this way on purpose. Once its (few, but powerful) features are internalized, it simply gets out of the way.

Read: *Rc – The Plan 9 Shell*, `rc(1)`

#### 8.1.1 – Prompts

Creating an `rc` function with the same name as your prompt allows you to easily double-click to select at the end of a previously typed line and then send it using the mouse button 2 menu (see the discussion of `rio` menus, below). This can be used to approximate a form of command history (see also the commands `"` and `"'`, which print and execute the previous command, respectively).

Add something like this to your `$home/profile`:

```
fn term%{ $* }
```

In `rc` the `;` character forces the end of a line and is treated as a noop when it appears alone, so it is also possible to create a simple prompt that would require no special prompt function in order for the prompt to be effectively ignored when selecting and sending:

```
prompt='; '
```

Obviously, the prompt can be named however the user sees fit.

### 8.1.2 – `/env`

**Note:** Contents of the `/env` directory are provided by the kernel and represent a separate accounting of the shell's environment; `rc` reads `/env` only on startup, and flushes/writes `/env` only before executing programs.

### 8.2 – `rio`



`rio` is the Plan 9 window system. More accurately, `rio` multiplexes input devices with and serves a file interface to a series of rectangles, inside the boundaries of which are drawn an arbitrary arrangement of pixels. Controlling the rectangles is more straightforward, and at the same time more flexible, than what is commonly expected from most "window managers."

Read: `rio(1)`, `rio(4)`

To effectively use `rio`, you need a three button mouse. If you only have a two button mouse you can emulate the middle button by holding down the `shift` key whilst pressing the right button.

**Note:** Button 1, 2, and 3 are used to refer to the left, middle, and right buttons respectively.

### 8.2.1 – The Pop-up Menu

Pressing and holding down mouse button 3 on the gray desktop or on a shell window will give you a menu with the following options:

New

Resize

Move

Delete

Hide

Pressing and holding down mouse button 2 results in a menu with the following options:

cut

paste

snarf

plumb

look

send

scroll

Select an item by releasing the button over the menu item. Rio uses the same button that started an action throughout that operation. If you press another button during the action the operation is aborted and any intermediate changes are reversed.

Each menu acts as a action verb selector which then requires an object (i.e. window) to be picked to indicate which window the verb is to act on. A further mouse action may then be required.

### 8.2.2 – Window control

Clicking on a window brings it to the front.

You can directly change the shape of a window by clicking and dragging on the edge or corner of the window border. Mouse button 1 or 2 will allow you to drag the edge or corner to a new size, and mouse button 3 will allow you to move the window.

The mouse button 3 menu contains a list of all windows that are currently obstructed by

other windows. Selecting a label tops the window.

The pop-up menu remembers the last command chosen, so as a shortcut you can just press and release button 3 without moving the mouse between pressing and releasing to select the previous command again.

In addition, `rio` serves a variety of files for reading, writing, and controlling windows. Some of them are virtual versions of system files for dealing with the display, keyboard, and mouse; others control operations of the window system itself. These files, as well as the `window(1)` command, allow for controlling windows programmatically by reading and writing text strings. Thus simplifying the automated opening and placement of various windows with user scripts.

Read: `rio(4)`

### 8.2.3 – Text in `rio` windows

Text in a `rio` window may be freely manipulated, edited, altered, deleted and/or acted upon using either mouse chords or the options from the mouse button menus. (For an example, see the discussion of the use of `rc` prompts, above.)

The special file `/dev/text` (for the current window), or `/dev/wsys/n/text` (for window `n`) contains all text that has already appeared in the window. The contents of this file may serve as a primitive form of command history (and may be acted upon using standard command line tools), but are lost when the window is closed.

Seriously, read: `rio(4)`

### 8.2.4 – Scrolling

By default, a `rio` window will fill up with text and then block.

Endless scrolling may be enabled by selecting `scroll` from the mouse button 2 menu.

The `up` or `down` arrow keys and `pgup` or `pgdwn` keys may be used to scroll up or down in consistently measured increments.

Holding down the `shift` key and pressing the up or down arrow key will scroll a single line in the respective direction.

9front's `rio` supports mousewheel scrolling. The heuristic employed is roughly the same as that of clicking in the scrollbar on the left of the window: when the mouse pointer is near the top of the window the scrolling increment is small, while as the mouse pointer approaches the bottom of the window the scrolling increment grows progressively larger. Presently this behavior is limited to `rio`, `sam`, and `mothra` but may later be extended to other programs.

**Note:** While the behavior of the arrow and page keys is fairly consistent between programs, mousewheel scrolling is not. So far, `shift + up` or `down` is only supported in `rio` windows.



## 8.2.5 – Mouse Chording

Almost anywhere — `sam(1)`, `acme(1)`, `window(1)` — you can use the following mouse chords:

`mb1` — Select text.

`mb1 double click` — Select word under cursor, or at the end/start of a line, select the whole line.

After selecting with `mb1` and while still holding `mb1` down (these chords also work with text selected by double-clicking, the double-click expansion happens when the second click starts, not when it ends):

`mb2` — Cut text.

`mb3` — Paste text (can be reverted by clicking `mb2` immediately afterwards).

To snarf (copy), click `mb2` immediately followed by `mb3`.

## 8.2.6 – Keyboard Shortcuts

Almost anywhere — `sam(1)`, `acme(1)`, `window(1)` — you can use the following shortcuts:

`Ctrl-u` — Delete from cursor to start of line.

`Ctrl-w` — Delete word before the cursor.

`Ctrl-h` — Delete character before the cursor.

`Ctrl-a` — Move cursor to start of the line.

`Ctrl-e` — Move cursor to end of the line.


`Ctrl-b` — Move cursor to the position immediately after the prompt. (`rio` only)

Read: UNIX Keyboard Bindings

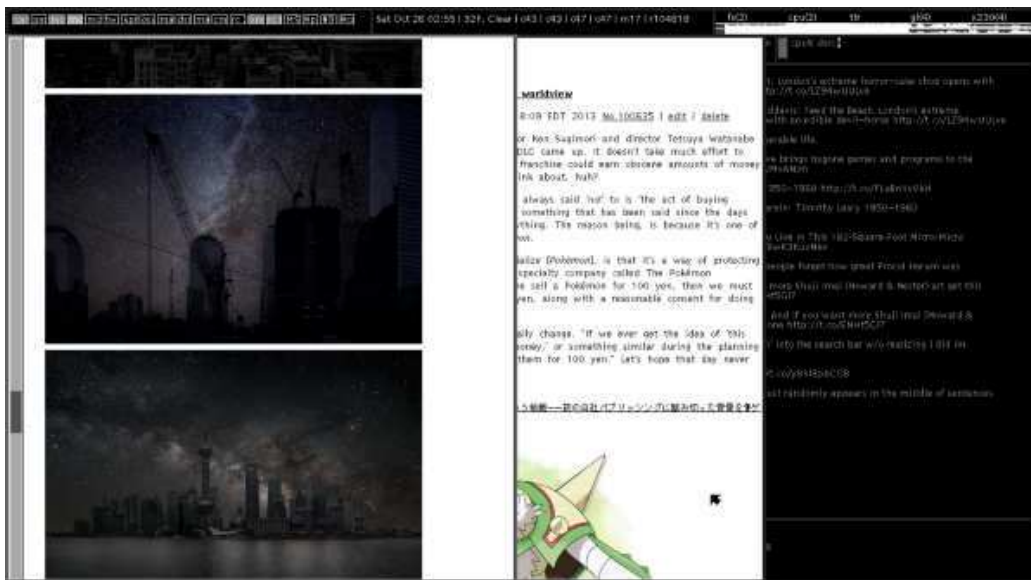
In a `rio(1)` window, scroll up or down one line by holding `shift` and pressing the up or down arrow.

## 8.2.7 – Color scheme

`rio` looks like this:



rio's color scheme may be modified by editing the .c configuration files and recompiling:



**Note:** Someone will mock you for doing this.

See: <http://plan9.stanleylieber.com/rio>

Rob Pike, rio's author, was all like:

the clean appearance of the screen comes mostly from laziness, but the color scheme is (obviously) deliberate. the intent was to build on an observation by edward tufta that the human system likes nature and nature is full of pale

colors, so something you're going to look at all day might best serve if it were also in relaxing shades. renee french helped me with the specifics of the color scheme (she's a professional illustrator and my color vision is suspect), once i'd figured out how i wanted it to look. there are still some features of the color system that i put in that i think no one has ever noticed. that's a good thing, in my opinion; the colors should fade away, if you'll pardon the expression. having used other systems with different approaches to color screens, most especially windows XP (extra pukey), i think tufte was right.


Rob Pike, 2003

The color scheme was an attempt to honor a point made originally in a little brochure by Edward Tufte that the colors of nature are soft and quiet and peaceful to look at, while most computer screens are covered in glaring bright colors. When color came to the system I wanted it to be pleasant.

Rob Pike, 2008

See: [edwardtufte.com](http://edwardtufte.com)

### 8.2.8 – Why is rio like this?



Window systems should be transparent. That's the argument put forward in the famous paper by rio's author, Rob Pike.

Beyond this, Rob offered an explanation (in response to a question on the 9fans mailing list) of some of the choices made in the design of 8½ and rio:

> functioning cursor keys would still be a speed benefit.

This feels true but is false. There were some fascinating experiments done a few years ago in which people were given a long, tedious editing task. Some

of the people were keyboard fans, some were mouse fans. Both folks were asked to do the task two ways, in random order, once using the mouse to do the editing, once using cursor keys etc. Regardless of their predilections, which was stated up front, after the experiment everyone who did the task agreed that it was faster to use the keyboard than the mouse to complete the task. Everyone. Here's the kicker: everyone was wrong. They were being timed, and in fact the reverse was true. Although they thought the keyboard was faster, doing the task using the mouse was faster for everyone, by a substantial fraction.

The explanation, besides the obvious that arrow keys are actually pretty slow if you're going more than a line or character, is that people feel the mouse wastes time because you need to grab it and move it, but it's time well spent. The part of the brain that uses keyboard commands to move the cursor is a higher-order function, and thinking and planning how to use the keys to get to the destination blocks thinking about the editing task at hand. But using the mouse is done by a lower-order part of the brain, which keeps the editing part of the brain clear. There's less task switching going on when you use the mouse, so you work more efficiently.

If you don't believe me, the story is here:

<http://www.asktog.com/readerMail/1999-12ReaderMail.html>

Thanks to some forgotten 9fan who mentioned this a while back. I didn't know about these experiments when I said, long ago, that using arrow keys to point at a display is like telling someone how to go somewhere by giving directions, while using a mouse is like pointing at a map. In fact, I never used a screen editor until I had a mouse, for just this reason.

Rob Pike, 2001

## 8.2.9 – tips

### 8.2.9.1 – Taking a screenshot

To capture the entire screen:

```
topng </dev/screen >screen.png
```

To capture only the current window:

```
topng </dev/window >window.png
```

It is also possible to capture *other* windows:

```
topng </dev/wsys/n/window >window.png
```

where *n* is the number of the window being captured.

Read: `rio(4)`

## 8.3 Text Editors

### 8.3.1 – sam

The text editor `sam` was created by Rob Pike for Research UNIX V9 (circa 1986), and later included with Plan 9.

See: <http://sam.cat-v.org>


Read:

*The Text Editor sam* — The original paper by Rob Pike.

*A Tutorial for the Sam Command Language* — Documents the editing language.

*sam quick reference card*

`sam(1)` man page



Dan Flavin, *Document for Untitled (to the “innovator” of Wheeling Peachblow)*, 1968

#### 8.3.1.1 – Scrolling

9front’s slightly modified version of `sam` supports mousewheel scrolling in the same manner as `rio`.

Read: *FQA 8.2.4 – Scrolling*

### 8.3.1.2 – Mouse Chording

9front sam supports the same mouse chording as `rio`.

Read: *FQA 8.2.5 – Mouse Chording*

### 8.3.1.3 – Why does sam have a separate snarf buffer from rio?

The program's author, Rob Pike, says:

was a consequence of running over 1200 baud when sam was first written. you didn't want every cut and paste to bounce off the remote end at that speed. nowadays that argument has less weight. on the other hand, i still kinda like that you can have an editing session that doesn't corrupt what you have in rio's snarf buffer. i tried the unified way in acme and i often (not always) miss the old way.

Rob Pike, 2003


### 8.3.1.4 – Keyboard Shortcuts

Esc — Cut (and consequently, snarf) the selected text.

Ctrl-b — Switch focus to the edit window.

### 8.3.2 – acme

*There is also an alternative user interface, acme(4), that some people use as their editor.*  
— Geoff Collyer



Handmade cyclogram by Russian cosmonaut, Georgi Grechko.

## 8.4 – Internet



Sending and receiving bits via alien protocols.

### 8.4.1 – Mail

Read: mail(1), FQA 7.7 – Mail server configuration and maintenance

### 8.4.1.1 – upasfs

From upasfs(4):

Fs is a user level file system that caches mailboxes and presents them as a file system. A user normally starts fs in his/her profile after starting plumber(4) and before starting a window system, such as rio(1) or acme(1). The file system is used by nedmail(1), acme(1)'s mail reader, and imap4d and pop3 (both pop3(8)) to parse messages. Fs also generates plumbing messages used by biff and faces(1) to provide mail announcements.

Read: upasfs(4), pop3(8), faces(1)

#### 8.4.1.1.1 – Reading gmail via IMAP

```
upas/fs -f /imaps/imap.gmail.com/your.username@gmail.com
```

The first time this command is run, you should see an error that looks something like this:

```
upas/fs imap: server certificate 22471E10D5C1E41768048EF5567B27F532F33
not recognized
upas/fs: opening mailbox: bad server certificate
```

To add this certificate to your system, type:

```
echo 'x509 sha1=22471E10D5C1E41768048EF5567B27F532F33' \
>>/sys/lib/tls/mail
```

Once upas/fs is running, you can open as many gmail mailboxes (labels) as you wish:

```
echo open /imaps/imap.gmail.com/your.username@gmail.com/yourlabel \
yourlabel >/mail/fs/ctl
```

**Note:** upas/fs reads the entire mailbox into RAM, which means it may fail to load mailboxes that exceed the available system memory. This problem has been addressed in Erik Quanstrom's nupas.

### 8.4.1.2 – nedmail

nedmail is a command line mail client similar to the classic mail client shipped with Research UNIX.

Read: nedmail(1)

### 8.4.1.3 – nupas

Read: *Scaling Upas*, by Erik Quanstrom

**Note:** Installing nupas replaces the system's default upas. Read the included README.9front and mkfile before attempting to install.



## 8.4.2 – NNTP

Read: `newt(1)`, `nntpfs(4)`

## 8.4.3 – IRC

### 8.4.3.1 – ircrc

`ircrc` is an IRC client implemented in `rc`. It is included with 9front.

Read: `ircrc(1)`

### 8.4.3.2 – irc7

A persistent IRC client was written in the `c` programming language by Andrey Mirtchovski. It has been modified slightly by 9front users (mainly, adding an `-e` flag to the `ircsrv` program that implements SSL connections).

## 8.4.4 – FTP

Read: `ftpfs(4)`

## 8.4.5 – HTTP



### 8.4.5.1 – mothra



`mothra` is a trivial web browser written in 1995 by Tom Duff. It ignores Javascript, CSS and many HTML tags. It was dropped from Plan 9 after the 2nd Edition, but has been picked up and (somewhat) refined for 9front.

Read: `mothra(1)`

### 8.4.5.2 – abaco

no.

### 8.4.5.3 – hget

`hget` is a command line HTTP client (similar to programs such as `curl` or `wget`) that started out as a c program in Plan 9 from Bell Labs, but was re-implemented in `rc` for 9front.

Read: `hget(1)`

### 8.4.5.4 – charon

The Inferno operating system can be run hosted on Plan 9, and includes a GUI web browser called `charon`, which implements ECMASCRIPT 1.0 as well as additional HTML attributes.

**Note:** `charon` is ancient and is not really a sufficient replacement for 9front's web browsers. The rudimentary javascript support can be useful for some simple tasks.

#### 8.4.5.5 – i

There exists an unfinished/buggy port of `charon` from Inferno's limbo programming language to Plan 9 c.

#### 8.4.6 – SSH

Several SSH clients exist for Plan 9, none of which are perfect.

##### 8.4.6.1 – ssh

9front ships with the original Plan 9 native SSH1 client from Bell Labs.

Read: `ssh(1)`

##### 8.4.6.2 – ssh2

Programmers at Coraid created a Plan 9 native SSH2 client that was picked up (and completely rewritten) by Bell Labs. It is currently not included with 9front.

**Note:** There are bugs and expected features are missing. Consult the source.

##### 8.4.6.3 – scpu

Two 9front users (taruti and mischief) worked on an SSH2 client written in the Go programming language. It has been extended to work with Plan 9 factotum(4), but still does not fully honor complex Plan 9 dial(2) strings.

##### 8.4.6.4 – OpenSSH

Plan 9 user fgb ported OpenSSH 4.7p1, OpenSSL 0.9.8g 19 Oct 2007 to Plan 9. It is available in his contrib directory, or a 386 binary is available here: [openssh.tgz](#).

##### 8.4.6.5 – sftpfs

An implementation of sftpfs was created for Plan 9 that can work with either the native SSH clients or fgb's OpenSSH port.

###### 8.4.6.5.1 – Mounting a remote u9fs share over SSH

The `u9fs` program runs on UNIX and serves an unencrypted 9P(2) share. It is possible to mount such a share over SSH.

With `ssh2`:

```
srv -s 5 -e 'ssh2 -l sl wm '/usr/local/bin/u9fs \  
-u sl -na none''' wm /n/wm
```

With `scpu`:

```
srv -s 5 -e 'scpu -u sl -h wm -c \  
''/usr/local/bin/u9fs -u sl -na none''' wm /n/wm
```

In both cases, an ssh connection is opened to remote UNIX host `wm`, logged in with user `sl` and mounted on Plan 9 under `/n/wm`.

Read: `u9fs(4)`, `srv(4)`

### 8.4.7 – `secstore`

Typing in lots of passwords over and over again is annoying.

`Secstore` authenticates to a `secure-store` server using a password and optionally a hardware token, then saves or retrieves a file. This is intended to be a credentials store (public/private keypairs, passwords, and other secrets) for a `factotum`.

Read: *FQA 7.4.3 – `secstored`* for information on setting up the `secstore` server, and: *FQA 7.4.3.1 – `Adding users to secstore`* to add users.

Finally, to add a secret to the file read by your `factotum` at startup, open a new window, type

```
ramfs -p; cd /tmp  
auth/secstore -g factotum  
secstore password:  
echo 'key proto=apop dom=x.com user=ehg !password=hi' >> factotum  
auth/secstore -p factotum  
secstore password:  
read -m factotum > /mnt/factotum/ctl
```

and delete the window. The first line creates an ephemeral memory-resident workspace, invisible to others and automatically removed when the window is deleted. The next three commands fetch the persistent copy of the secrets, append a new secret, and save the updated file back to `secstore`. The final command loads the new secret into the running `factotum`.

The `ipso` command packages this sequence into a convenient script to simplify editing of files stored on a secure store. It copies the named files into a local `ramfs` and invokes `ditor` on them. When the editor exits, `ipso` prompts the user to confirm copying modified or newly created files back to `secstore`. If no file is mentioned, `ipso` grabs all the user's files from `secstore` for editing.

By default, `ipso` will edit the `secstore` files and, if one of them is named `factotum`, flush current keys from `factotum` and load the new ones from the file.

Read: `secstore(1)`, `secstore(8)`

### 8.5 – `Audio`

Read: `audio(1)`, `audio(3)`

## 8.6 – External Media

### 8.6.1 – Mount an ISO9660 CD-ROM

```
mount <{9660srv -s} /n/iso /dev/sdD1/data # cd-rom drive
```

or:

```
mount <{9660srv -s} /n/iso /path/to/9front.iso
```

Read: dosrv(4)

### 8.6.2 – Burn a CD-ROM

```
cdfs  
cp 9front.iso /mnt/cd/wd  
rm /mnt/cd/wd
```

Read: cdfs(4)

### 8.6.3 – Mount a FAT formatted USB device

FAT formatted USB devices are automatically mounted under the `/shr` directory.

**Note:** Devices must be FAT or FAT32 formatted; exFAT is not supported.

## 8.7 – Emulation

### 8.7.1 – Linux Emulation



`linuxemu` is a program that can execute Linux/i386 ELF binaries on Plan 9. Semi-modern web browsers and other Linux programs may be run using `linuxemu` (if necessary, in conjunction with the `equis` X11 server).

**Note:** `linuxemu` can only be run on a Plan 9 system booted with a 386 kernel and binaries.

## BOOTSTRAP

To run `linuxemu`, you need a Linux root file system packed into a tarball:

[http://9hal.ath.cx/usr/cinap\\_lenrek/mroot-linuxemu.tbz](http://9hal.ath.cx/usr/cinap_lenrek/mroot-linuxemu.tbz)

<http://plan9.stanleylieber.com/linuxemu/mroot.tgz>

The `mroot-linuxemu.tbz` version contains no symlinks and can be extracted with plain Plan 9 tools `bunzip` and `tar`.

The `mroot.tgz` version contains the same Debian Sarge base as `mroot-linuxemu.tbz`, but with several additional packages pre-installed:

`9base`

`dmenu-4.1.1`

`dwm-5.8.2`

`gcc 3.3.5`

`linux-kernel-headers`

`mercurial 0.9.4`

`opera 10.11`

`python 2.3.5`

`xlib-dev`

and more.

You can create your own `mroot` with `debootstrap` on Debian Linux, or help write an installer that unpacks and installs an alternative distribution on Plan 9... In any case, `linuxemu` is not hardwired to any Linux distribution!

## RUNNING

Use the provided `linux` script to chroot into your Linux `mroot`. The `linux` script is necessary because for Linux programs to run, shared libraries from your `mroot` have to appear in its `/lib` and `/usr/lib` directories, while configuration files are expected to be in `/etc`. The script will build a private namespace and then bind the Linux `mroot` over the Plan 9 root. The original Plan 9 namespace is mounted within `linuxemu` under `/9`.

Assuming `mroot` is located in the current directory, start `linuxemu` like this:

```
linux -r ./mroot /bin/bash -i
```

If the `-r` option is omitted, the Linux `mroot` defaults to `/sys/lib/linux` on the Plan 9 machine.

In the Linux `mroot`, `/etc/resolv.conf` should be changed to match your network nameserver. In addition, `/etc/apt/sources.list` should be updated to a working Debian mirror. Sarge packages can still be accessed at:

```
deb http://archive.debian.org/debian-archive/debian sarge
main
```

## EXAMPLES

Linux X11 programs may be used in conjunction with the `equis` X11 server. For example, to run the Opera web browser under your Linux `mroot`, start `equis` in a `rio` window, start `linuxemu` in another `rio` window and then from within `linuxemu`:

```
dwm & # X11 window manager
opera & # web browser
```

Opera should (eventually) appear in the `equis` window. A window manager (this example uses `dwm`) is recommended so that X11 programs interact with window resources properly.

## DEBUGGING

If `linuxemu` crashes, use `acid` to figure out what's going on:

```
mk acid
acid -l linuxemu.acid <pid>
```

Then you can issue the following commands:

```
ustk()
```

dump a (userspace) stacktrace for the current thread:

```
umem(Current())      dump the memory mappings
ufds(Current())      dump the filedescriptor table
ustrace(Current())   dump the internal tracebuffer
                     (enabled by -d option)
```

Use `xasm()` and `xcasm()` for disassembly for Linux code.

Read: `acid(1)`

You can also enable full trace logging:

```
linux -r ./mroot -dd /bin/bash -i >[2]/tmp/linuxemu.log
```

This slows `linuxemu` down considerably. In case of race conditions, it often happens that the bug disappears when doing full trace logging!

## 8.7.2 – Nintendo



Emulators for several Nintendo video game consoles ship with the system:

gb — Game Boy

gba — Game Boy Advance

nes — Nintendo Entertainment System

snes — Super Nintendo Entertainment System

Read: `nintendo(1)`



### 8.7.3 – Sega



An emulator for the Sega Megadrive/Genesis video game console ships with the system:

md — Sega Mega Drive/Genesis

Read: [sega\(1\)](#)

### 8.7.4 – Commodore



An emulator for the Commodore 64 home computer ships with the system:

c64 — Commodore 64

Read: `commodore(1)`

## 8.8 – Additional Software

### 8.8.1 – 9front sources server

Additional 9front software is available from a 9P share that is accessible from any Plan 9 system:

```
9fs 9front.org
```

The following files and directories will then be available under `/n/9front.org/`:

```
9front/ — 9front source
```

```
9front.torrent — torrent of current 9front ISO image
```

```
extra/ — third party software source
```

```
iso/ — current 9front ISO image(s)
```

```
lists/ — mailing list archives)
```

```
rc/ — third party rc scripts
```

### 8.8.2 – 9front contrib

Some 9front users maintain a contrib directory on an official 9front 9P share (similar to the contrib arrangement provided by Bell Labs) that is accessible from any Plan 9 system:

```
9fs contrib.9front.org
```

The following files and directories will then be available under `/n/contrib.9front.org/`:

```
contrib/ — directory tree containing user directories
```

```
lsr — index of files under contrib/ (generated daily)
```

## 8.9 – Bootstrapping architectures not included on the ISO

### 8.9.1 – amd64

To setup the amd64 port, install the 386 port from the ISO, then cross compile and install the amd64 binaries and kernel.

Read: *FQA 5.2.2.1 – Cross compiling*, *FQA 7.2.5 – How do I install a new kernel?*


## 8.9.2 – Raspberry Pi

The most convenient way to use an rpi is to cross compile and install the `arm` binaries and the `bcm` kernel on the network file server, and then `tcp` boot the rpi.

Read: *FQA 5.2.2.1 – Cross compiling, FQA 6.7.1 – How do I tcp boot?*

Outdated and possibly confusing instructions for installing directly onto the rpi's sd card are detailed in *Appendix J – Junk*

## 8.10 – ACPI



Plan9front currently has partial ACPI support for PCI interrupt routing and system shut-down.

### 8.10.1 – Enabling ACPI

ACPI is enabled with the presence of `*acpi=` boot parameter.

This will create the `/dev/acpitb1s` file that can be used to read the systems acpi tables. Specifying `*acpi=0` will make acpi tables accessible thru the file but not use it in the kernel.

### 8.11 – MP Table Override

*This horrible hack should be considered obsolete, in particular with the implementation of APIC renaming and MSI. It might still be useful.*

Since some motherboard/BIOS manufacturer don't care about anything but ACPI, the MP table might be in an unusable state which meant that `*nomp=1` had to be used. The new MP Table Override can be used to specify an entirely different MP table for 9front to

use.

To specify an MP table the `*mp` and `*mpN` options are used, for example (for illustration purposes, not a valid mp table):

```
*mp=16
*mp0=00 11 33 13 15
*mp1=16 99 DD 1C
*mp2=9F 81 2a 0a 7b 8c 1c
```

`*mp` specifies the size in byte and `*mpNR`, where *N* is a consecutive number starting with zero, specifies the data in hex form.

The in-hardware table can be dumped with `*dumpmp=1` (PROTIP: use a serial console).

`/sys/src/9/pc/mp.h` is useful to have handy.

## To Do

(Some are mutually exclusive)

- Write an utility to generate a proper table from Linux or similar
- Write user space tools to manipulate the mp table
- Calculate size automatically (and drop `*mp`)
- Use a text based format parsed by the kernel
- Fix MP support altogether (might not be feasible)

## Example

For the HP Compaq DC5700:

```

*mp=528
*mp0=00 00 10 03 65 0f 00 00 ff fb eb bf 00 00 00 00
*mp1=00 00 00 00 00 01 10 01 65 0f 00 00 ff fb eb bf
*mp2=00 00 00 00 00 00 00 00 01 00 50 43 49 20 20 20
*mp3=01 01 50 43 49 20 20 20 01 07 50 43 49 20 20 20
*mp4=01 20 50 43 49 20 20 20 01 ff 49 53 41 20 20 20
*mp5=02 02 11 01 00 00 c0 fe 03 00 0f 00 00 08 02 10
*mp6=03 00 0f 00 00 09 02 11 03 00 0f 00 00 0c 02 10
*mp7=03 00 0f 00 00 0d 02 11 03 00 0f 00 00 0e 02 12
*mp8=03 00 0f 00 00 74 02 14 03 00 0f 00 00 75 02 15
*mp9=03 00 0f 00 00 76 02 16 03 00 0f 00 00 77 02 17
*mp10=03 00 0f 00 00 68 02 14 03 00 0f 00 00 69 02 15
*mp11=03 00 0f 00 00 6a 02 16 03 00 0f 00 00 01 02 13
*mp12=03 00 0f 00 00 6c 02 15 03 00 0f 00 00 7c 02 11
*mp13=03 00 0f 00 00 7d 02 12 03 00 0f 00 00 7e 02 12
*mp14=03 00 0f 00 00 70 02 14 03 00 0f 00 00 71 02 15
*mp15=03 00 0f 00 00 72 02 16 03 00 0f 00 00 73 02 17
*mp16=03 00 0f 00 07 10 02 14 03 00 0f 00 07 11 02 15
*mp17=03 00 0f 00 07 12 02 16 03 00 0f 00 07 13 02 17
*mp18=03 00 0f 00 07 24 02 15 03 00 0f 00 07 25 02 16
*mp19=03 00 0f 00 07 26 02 17 03 00 0f 00 07 27 02 14
*mp20=03 00 0f 00 07 2c 02 16 03 00 0f 00 07 2d 02 17
*mp21=03 00 0f 00 07 2e 02 14 03 00 0f 00 07 2f 02 15
*mp22=03 00 0f 00 07 34 02 14 03 00 0f 00 07 35 02 15
*mp23=03 00 0f 00 07 36 02 16 03 00 0f 00 07 37 02 17
*mp24=03 00 0f 00 20 00 02 10 03 00 0f 00 20 01 02 11
*mp25=03 00 0f 00 20 02 02 12 03 00 0f 00 20 03 02 13
*mp26=03 00 00 00 ff 01 02 01 03 00 00 00 ff 03 02 03
*mp27=03 00 00 00 ff 04 02 04 03 00 00 00 ff 06 02 06
*mp28=03 00 00 00 ff 07 02 07 03 00 00 00 ff 08 02 08
*mp29=03 00 00 00 ff 09 02 09 03 00 00 00 ff 0c 02 0c
*mp30=03 00 00 00 ff 0d 02 0d 03 00 00 00 ff 0e 02 0e
*mp31=03 00 00 00 ff 0f 02 0f 04 01 00 00 ff 00 ff 01
*mp32=01 3f 50 43 49 20 20 20 03 00 0f 00 3f 00 02 11

```

## 8.12 – Revision Control

### 8.12.1 – cvs

OpenCVS was ported to Plan 9.

An implementation of a cvs file server, called `cvsfs`, was also created for Plan 9.

### 8.12.2 – git

More than one person has claimed to have ported `git` to Plan 9, but no code has ever been made public. Please stop telling unsuspecting newbs that `git` is available for Plan 9.

### **8.12.3 – Mercurial**


9front ships with Mercurial.

Read: *FQA 5.2.1.1 – hgrc*


### **8.12.4 – svn**

No.

## FQA 9 – Troubleshooting



### 9.1 – First



- STATE YOUR ASSUMPTIONS.
- `cat /etc/os-release` to verify you are not, in fact, running Ubuntu Linux with a Plan 9 theme.
- Are you running `drawterm`?
- Verify your OpenBSD configuration.

- Are you using the qwerty keyboard layout?
- Before reporting a bug, try the latest ISO image.
- MAKE SENSE.
- When all else fails, see: *FQA 2.4 – Reporting Bugs*



## 9.2 – Booting

### 9.2.1 – Boot parameters

- Immediately after the BIOS screen, hit any key until you see the > prompt. From there, values from `plan9.ini` may be temporarily added or changed. Read: `9boot(8)`
- Boot parameters beginning with a \* are interpreted by the kernel. All other parameters are passed as environment variables to the boot process.
- Adding or changing a parameter: `param=value`
- Removing a parameter: `clear param=`
- When finished, type `boot` to resume booting.

### 9.2.2 – Break into a shell

At the `[bootargs]` prompt, type `!rc` and hit enter to break into a shell. Type `exit` to return to the `[bootargs]` prompt.



### 9.2.3 – Editing plan9.ini

- It is not possible to edit plan9.ini stored on the ISO, but parameters can be changed temporarily before booting. See above.
- On a harddrive installation, plan9.ini is stored with the bootloader and the kernel in a small FAT partition called 9fat at the beginning of the plan9 partition. The 9fat can be mounted by executing 9fs 9fat from the livecd or the installed system. The file /n/9fat/plan9.ini can then be edited with a text editor like acme(1), sam(1) or ed(1).
- If your change to plan9.ini or the 9fat made the system unbootable, it is always possible to manually override parameters on the > prompt (see above) or start the system from the livecd and pick your installed cwfs or hjfs partition on the [bootargs] prompt. Example: local!/dev/sdC0/fscache

### 9.2.4 – Boot media not recognized

- Break into a shell, then type `grep -n '^01' '#$/pci/'*ctl` to get the pci vid/did of the installed disk controllers. then look in the sd drivers to see if the controller is already recognized.
- Sometimes, there is a problem with the drive, not the controller. `cat /dev/sd*/ctl` to get the status of the individual drives.
- Try different BIOS settings like AHCI/IDE mode.
- Try the USB troubleshooting steps (see below). Sometimes USB problems prevent unrelated devices from working properly.

### 9.2.5 – I moved my hard drive between ports

If your file system is cwfs(4):

- Let's assume it went from sdE0 to sdE1.
- At the bootargs prompt:

```
local!/dev/sdE1/fscache -c
filsys main c(/dev/sdE1/fscache)(/dev/sdE1/fsworm)
filsys dump o
filsys other /dev/sdE1/other
end
```

## 9.3 – Graphics

### 9.3.1 – Rio fails to start

See the discussion of graphics in *FQA 4 – 9front Installation Guide*. When you find a working mode, update your `plan9.ini`.

## 9.4 – Networking

### 9.4.1 – Networking is not working

Read: *FQA 6.2.8 – Verifying network settings*

### 9.4.2 – Cannot resolve domain names

If `ndb/dns` is running but you are still unable to resolve domains, you can try adding a DNS server directly to `/net/ndb`.

Read: *FQA 6.2.5 – DNS Resolution*

### 9.4.3 – `/mnt/web/clone` does not exist

Programs that require `webfs(4)` to be running require `webfs(4)` to be running.

Read: `hget(1)`, `mothra(1)`

### 9.4.4 – PCMCIA WiFi stopped working after reboot

Do `fshalt` and power down completely instead of just rebooting with `fshalt -r`.

## 9.5 – USB

### 9.5.1 – Devices not recognized or not working

- Break into a shell, or simply type in a `rio` window: `cat /dev/usbevent` and try plugging in a USB device. If the devices were detected, some output should appear on screen. This doesn't mean we have a working driver for it but verifies that the USB controller and HUB driver recognized the device.
- Dump USB controller status to the console with `echo dump >/dev/usb/ctl` after that, you might recover the output from `/dev/kmesg` or run `cat /dev/kprint` in a separate `rio` window to not spill it all over the screen. Check the `nintr` and `tdintr` counters to see if they are all zero. If that's the case, then this might be an interrupt routing problem sometimes caused by broken/incomplete BIOS MP tables. See below.
- If USB keyboard is the only option, try to enable PS2 emulation in the BIOS and disable kernel usb support. See below.





*Every joke is a tiny revolution*  
— George Orwell

## Appendix J – Junk

### Raspberry Pi



#### Prebuilt image

sdcard image available for the raspberry pi:

[http://felloff.net/usr/cinap\\_lenrek/9front-pi-3125.9fcca5b8ac11.img.bz2.torrent](http://felloff.net/usr/cinap_lenrek/9front-pi-3125.9fcca5b8ac11.img.bz2.torrent)

<http://intma.in/downloads/9front-pi-3125.9fcca5b8ac11.img.bz2>

I used Richard Miller's `pi.uboot.sd.img` as a base, you may want to update the firmware on the FAT partition with files from here:

<https://github.com/raspberrypi/firmware>

older versions:

9front 95b0756087eb from Sun Oct 20 22:38:35 CES 2013

[http://intma.in/downloads/9front\\_pi.img.gz.torrent](http://intma.in/downloads/9front_pi.img.gz.torrent)

[http://intma.in/downloads/9front\\_pi.img.gz](http://intma.in/downloads/9front_pi.img.gz)

#### Installation Instructions

If you don't want to use the provided image, follow these instructions to make your own. You don't need a computer capable of running 9front, qemu will suffice.

## Image

Make a FAT partition containing all the necessary files to boot the Raspberry Pi. An easy way to do this is to use Richard Miller's `pi.uboot.sd.img` (it's small). Write it to your SD card or do the following if you're using qemu (installing to an SD card directly from qemu was **very** slow for me):

```
# set count to the size you want your image to be
# (in 512b units, see dd manpage)
dd if=/dev/zero of=9front-pi.img count=3782016

# copy minimal image to beginning of the newly created image
dd if=pi.uboot.sd.img of=9front-pi.img conv=notrunc
```

## Boot and Installation

Now boot from the 9front CD. For qemu see FQA 3.3.1 - qemu, this is what I did:

```
qemu-system-i386 \
    -enable-kvm \
    -drive file=9front-pi.img,if=virtio \
    -cdrom 9front-3125.9fcca5b8ac11.iso \
    -boot d
```

Now do the installation (see FQA 4 - 9front Installation Guide for help), use `hjfs`. Since you don't need a 9fat partition, you may want to delete the automatically created partitions and create your own (start at sector 2 or `prepdisk` will complain). Note that since you've removed 9fat, the installer will prompt you to do the `prepdisk` step again, just type the next step in manually.

After the installation you have a 9front system that only has 386 binaries. The next step is compile the arm binaries.

## Compiling for ARM

Reboot and boot from CD again, but this time use the newly created partition as root (e.g. `local!/dev/sdF0/fs`)

Before compiling you need to create some directories under `/arm/bin`:

```
echo allow >>/srv/hjfs.cmd      # allow changing ownership of files
cd /arm/bin
for(i in audio auth aux bitsy btc dial disk fs games ip ip/httpd ndb nusb pkg re
    mkdir $i
    chgrp -u sys $i
    chgrp sys $i
}
```

Then compile everything:

```
cd /sys/src
objtype=arm mk install
cd 9/bcm
mk install && mk install 'CONF=picpuf'
```

And copy the kernel to the FAT partition:

```
dossrv >/dev/null >[2]/dev/null
mount -c /srv/dos /n/dos /dev/sdF0/dos
cp /arm/9pif /arm/9picpuf /n/dos
```

Then set the kernel in `/n/dos/config.txt` to `9pif` or `9picpuf` for terminal or cpu kernel and boot the image with the Raspberry Pi. At this point you can delete the 386 binaries if you don't need them.

Once booted on the Raspberry Pi, you're able to mount the dos partition with:

```
9fs pidos
```

This will mount the dos partition on `/n/pidos`, similar to how `9fs 9fat` works.

## FQA Appendix L – Languages

C



C++





C#



Go




Read: <http://9.offblast.org/9/guide/go>

## Haskell



## Perl



Perl 5.8.0, ported to Plan 9: <http://plan9.bell-labs.com/sources/extra/perl.iso.bz2>

**PHP**



**Python**



Python 2.5.1 is included with the 9front distribution, not because anyone loves Python, but because it was required by Mercurial (also loved by no one), which was required by Google Code (shutting down in 2015). An abject lesson in expediency.



hjfs: ending

