


# An ISP Perspective, jail(8) Virtual Private Servers



Isaac Levy (.ike) <[ike@lesmuug.org](mailto:ike@lesmuug.org)>

Materials prepared for AsiaBSDCon 2007 Proceedings, University of Tokyo, Japan.  
These materials are Copyright 2006 Isaac Levy, under the terms of the BSD license.

# .ike Context

- I have lived inside, and outside, jails extensively for web application servers and software development purposes
- I am not a jail author, no commit bit...
- NO religious OS/Virtualization bashing, yet strong arguments presented based on specific application/threat models.

# .ike Context

- I helped run the ISP iMeme (now defunct), perhaps the first specific jail(8) based ISP in the US.
- I currently operate a micro, high-availability datacenter operation for my clients, using jail-oriented systems.

**(wintermute)** Jon Ringuette,  
(partner and founder of iMeme),  
taught me to jail(8).




**iMeme.net**  
"the infectious meme."



### Company Information

### Services and Technologies



- [solutions](#)
- [company](#)
- [client area](#)
- [search](#)
- [home](#)

### Welcome to iMeme

iMeme is a hosting provider dedicated to Open Source. We offer solutions and conveniences for web developers. We specialize in Zope and PHP and are building a hosted developers community. Each iMeme account is overseen directly by an owner of the company. We are experienced and kind administrators and provide virtual and dedicated server colocation with prompt, professional service.

### What We Offer

Our hosting packages provide you with the functionality and stability to make your ideas come to life. We specialize in reliability, systems and network administration, FreeBSD, and Zope.

iMeme offers a substantial knowledge base. Among our community of clients, we have some of the finest and most creative Zope developers. The iMeme knowledge base is focused in three ways. Canonical documentation and how-tos produced by it's managing owners (for iMeme to instruct clients), interactive documentation in the form of a wiki web (so that clients may teach iMeme), and real-time tools for assistance and support (such as irc.imeme.net).

We would like to have you join our growing community of developers and allow us to

**iMeme.net**  
"smart people getting along"

[news](#)

#### quick links:

[webmail](#)

[email password](#)

[policies](#)

[documentation](#)

[signup](#)

[contact us](#)

[mailing list](#)

[list archives](#)

Zope
Dynamic Functionality
Products
Email Services
IMAP Email Access
WebMail
Mailing List Manager
Domain Services
iMeme Rates
About Our Company
Contact Information
Documentation
Use Client Services
News
Sign up
email us

the control panel, install different  
ns of Zope at will. We provide you  
n Zope Enterprise Objects,

, administrate user accounts,

ope products, write external scripts,

resses at yours. We offer IMAP

# done technical talks on jail before *you have root!*

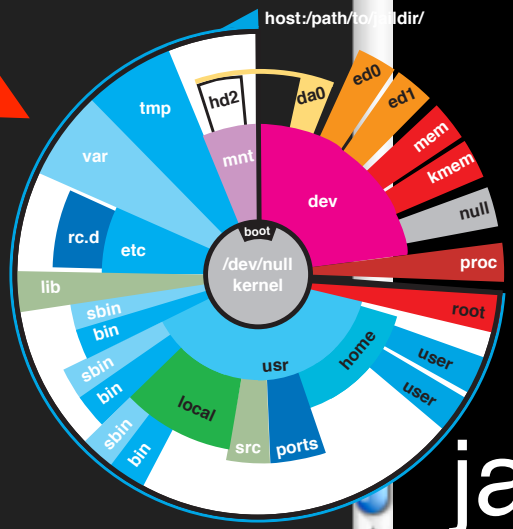
```
Terminal — ssh — 80x24
$ hostname
chick.diversaform.net
$ ifconfig
bge0: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
      options=1b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING>
      inet 10.0.1.200 netmask 0xffffffff broadcast 10.0.1.200
      ether 00:e0:81:34:bf:8c
      media: Ethernet autoselect (1000baseTX <full-duplex>)
      status: active
bge1: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> mtu 1500
      options=1b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING>
      ether 00:e0:81:34:bf:8c
      media: Ethernet autoselect (1000baseTX <full-duplex>)
      status: active
plip0: flags=10881<UP,BROADCAST,LOOPBACK,NOARP,POINTOPOINT> mtu 500
      options=1b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING>
      inet 10.0.1.200 netmask 0xffffffff broadcast 10.0.1.200
      ether 00:00:00:00:00:00
      media: Ethernet autoselect (1000baseTX <full-duplex>)
      status: active
pflog0: flags=0<>
pfsync0: flags=0<>
lo0: flags=8049<UP,LOOPBACK,RUNNING> mtu 1000
      options=1b<RXCSUM,TXCSUM,VLAN_MTU,VLAN_HWTAGGING>
      inet 127.0.0.1 netmask 0xff000000 broadcast 127.0.0.1
      ether 00:00:00:00:00:00
      media: Ethernet autoselect (1000baseTX <full-duplex>)
      status: active
$ su
Password:
chick# whoami
root
chick#
```



```
Terminal — ssh — 80x24
chickenhawk:/home/ike ike$ sudo cat /proc/35721/status
sshd 35721 1 35721 35721 - sldr 1138578391,270482 4,956102 27,766581 select 0 0
0,0 chick.diversaform.net
chickenhawk:/home/ike ike$ sudo cat /proc/93030/status
sshd 93030 93025 93025 93025 - no flags 1152553319,908960 0,61426 0,40950 select
1001 1001 1001,1001,1001,0 -
chickenhawk:/home/ike ike$
```



host



jail

# Warranty / Announcement

- This is a lot of information for 1 hour, will move fast.... (questions after, please)
- I'll be around if anyone has more complex questions or strategies they want to discuss.

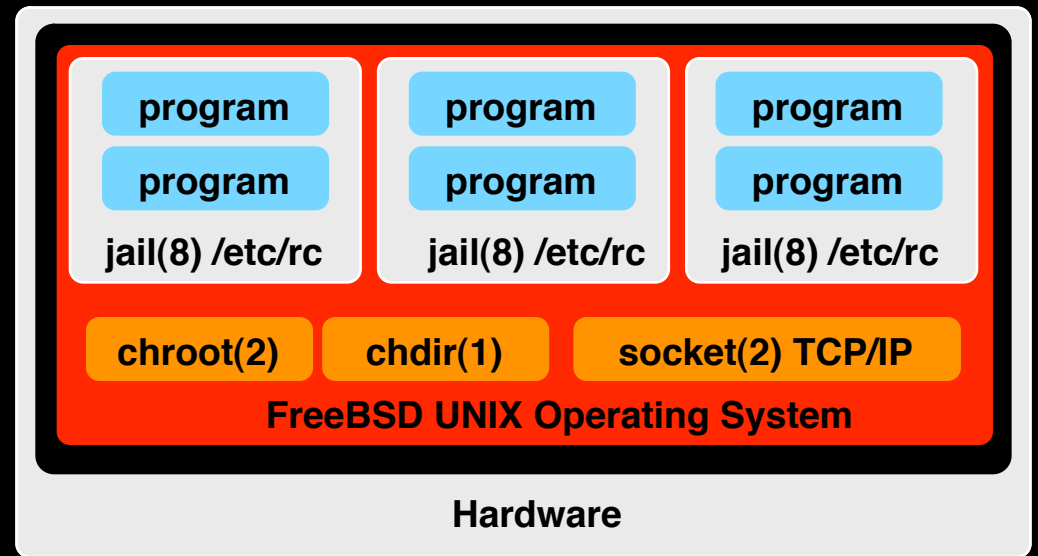

# Warranty / Announcement

- I'm *\*trying\** to stay close to classic UNIX process and ideas, and 'stock' methodology.
- I assume you all know the basics of the jail(8) mechanism on FreeBSD, but...




# jail(8) is:

**chroot(2) bound to an IP address**




# jail(8) is:



**chroot(2) bound to an IP address**

a tool for creating virtual servers

# jail(8) is:




**chroot(2) bound to an IP address**

a tool for creating virtual servers

native on FreeBSD since 1998

# jail(8) is:



**chroot(2)** bound to an **IP address**

a tool for creating virtual servers

native on FreeBSD since 1998

designed to partition  
**'mutually untrusted users'**

# Mutually Untrusted Users?



# Mutually Untrusted Users?




# Internet Service Provider (ISP)

**Common Definition:** a business or organization that provides users access to the Internet and related services, (web hosting, email, etc...).

**Abstract Definition:** providing users the ability to run programs, and maintain persistent data storage, available across a network (like the internet).

scaling, patterns, time  
(an exercise)






THE ARPA NETWORK

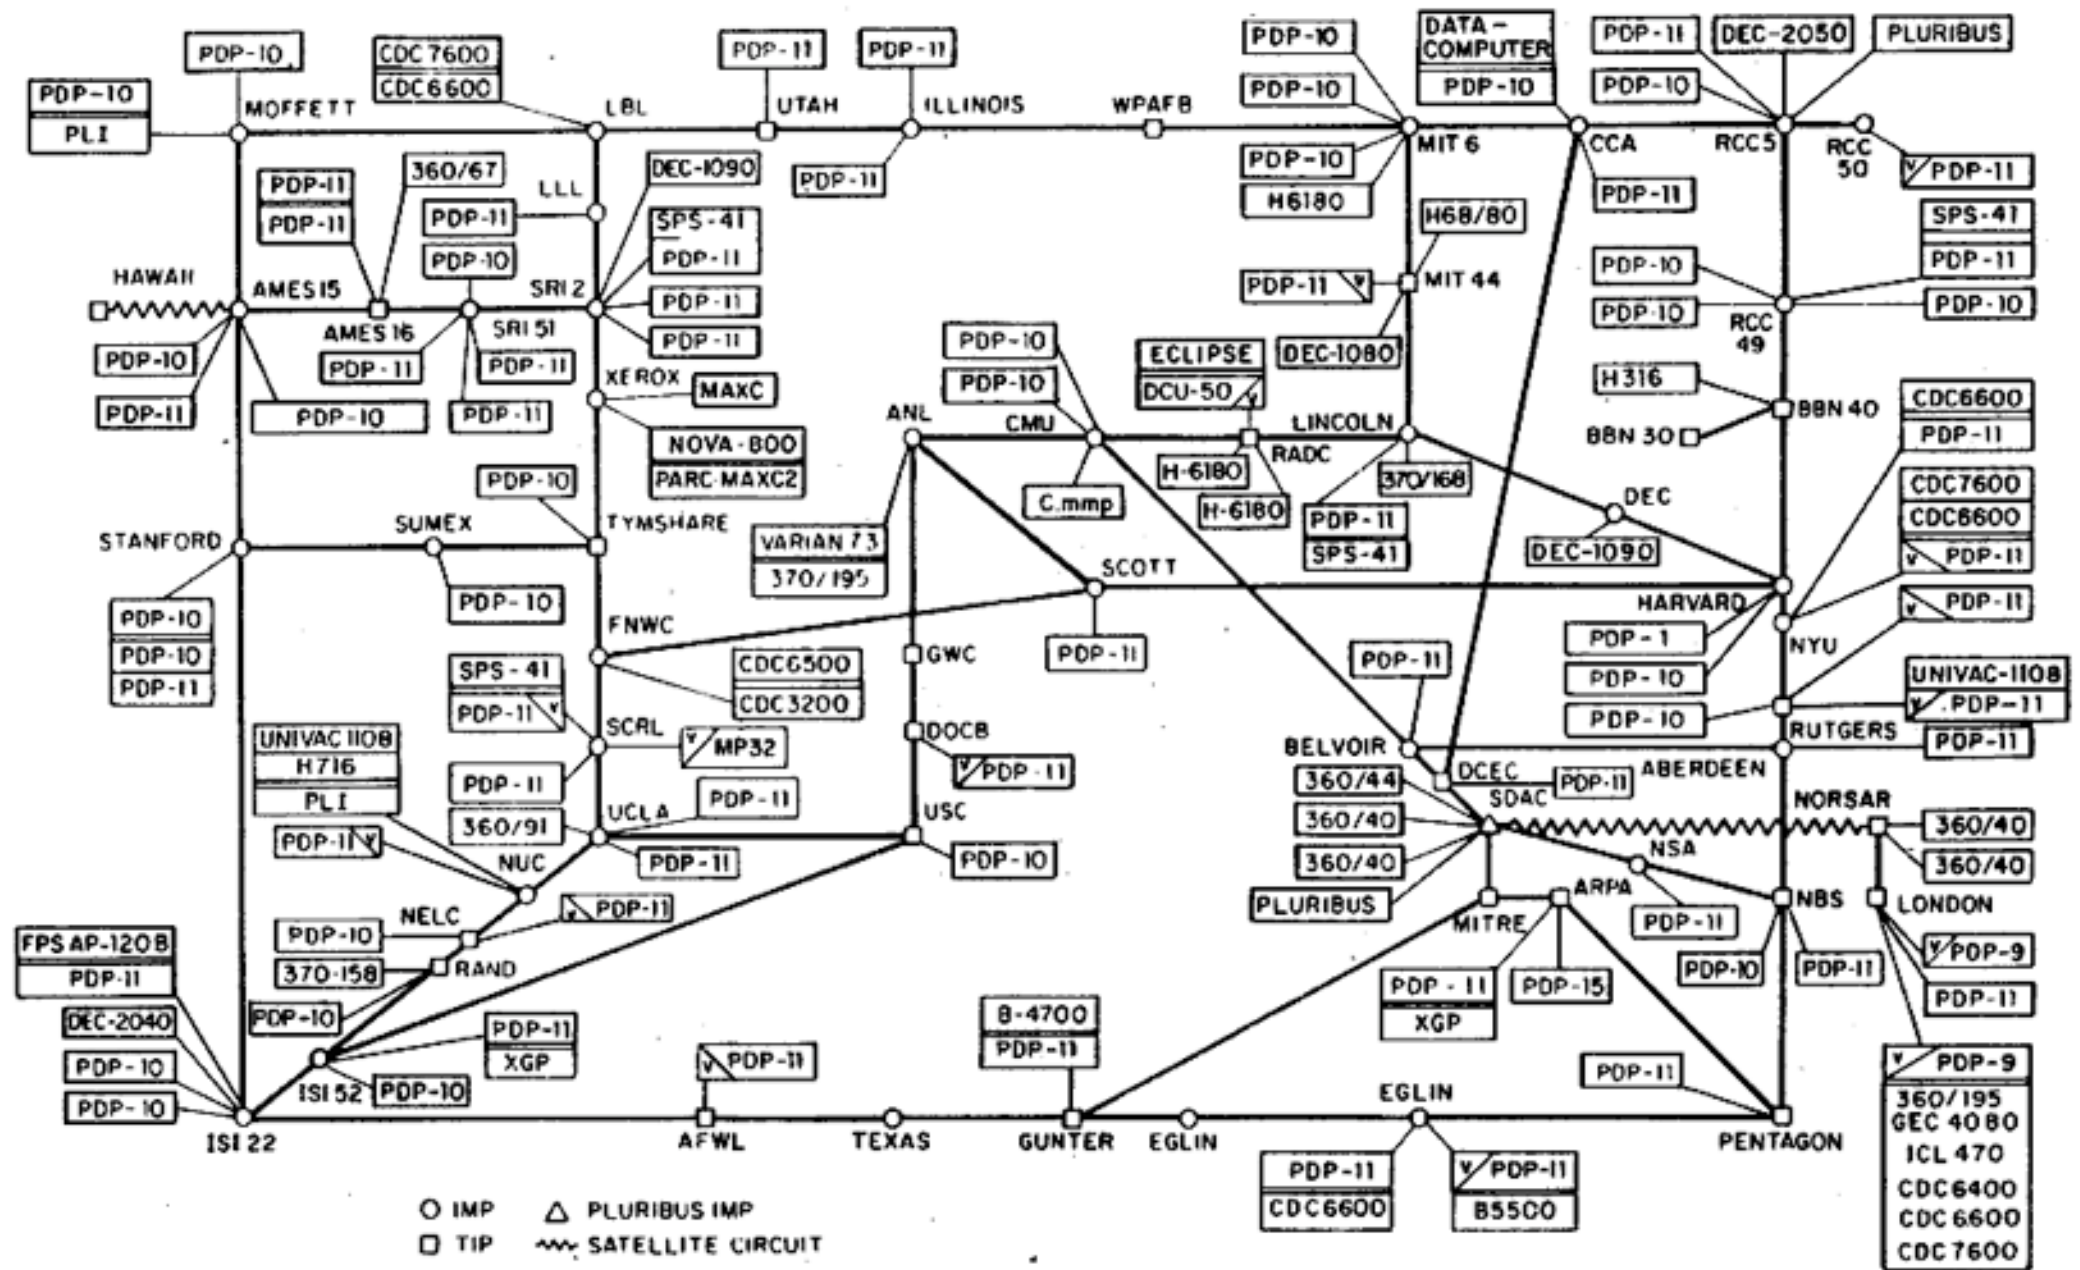
DEC 1969

4 NODES

FIGURE 6.2 Drawing of 4 Node Network  
(Courtesy of Alex McKenzie)



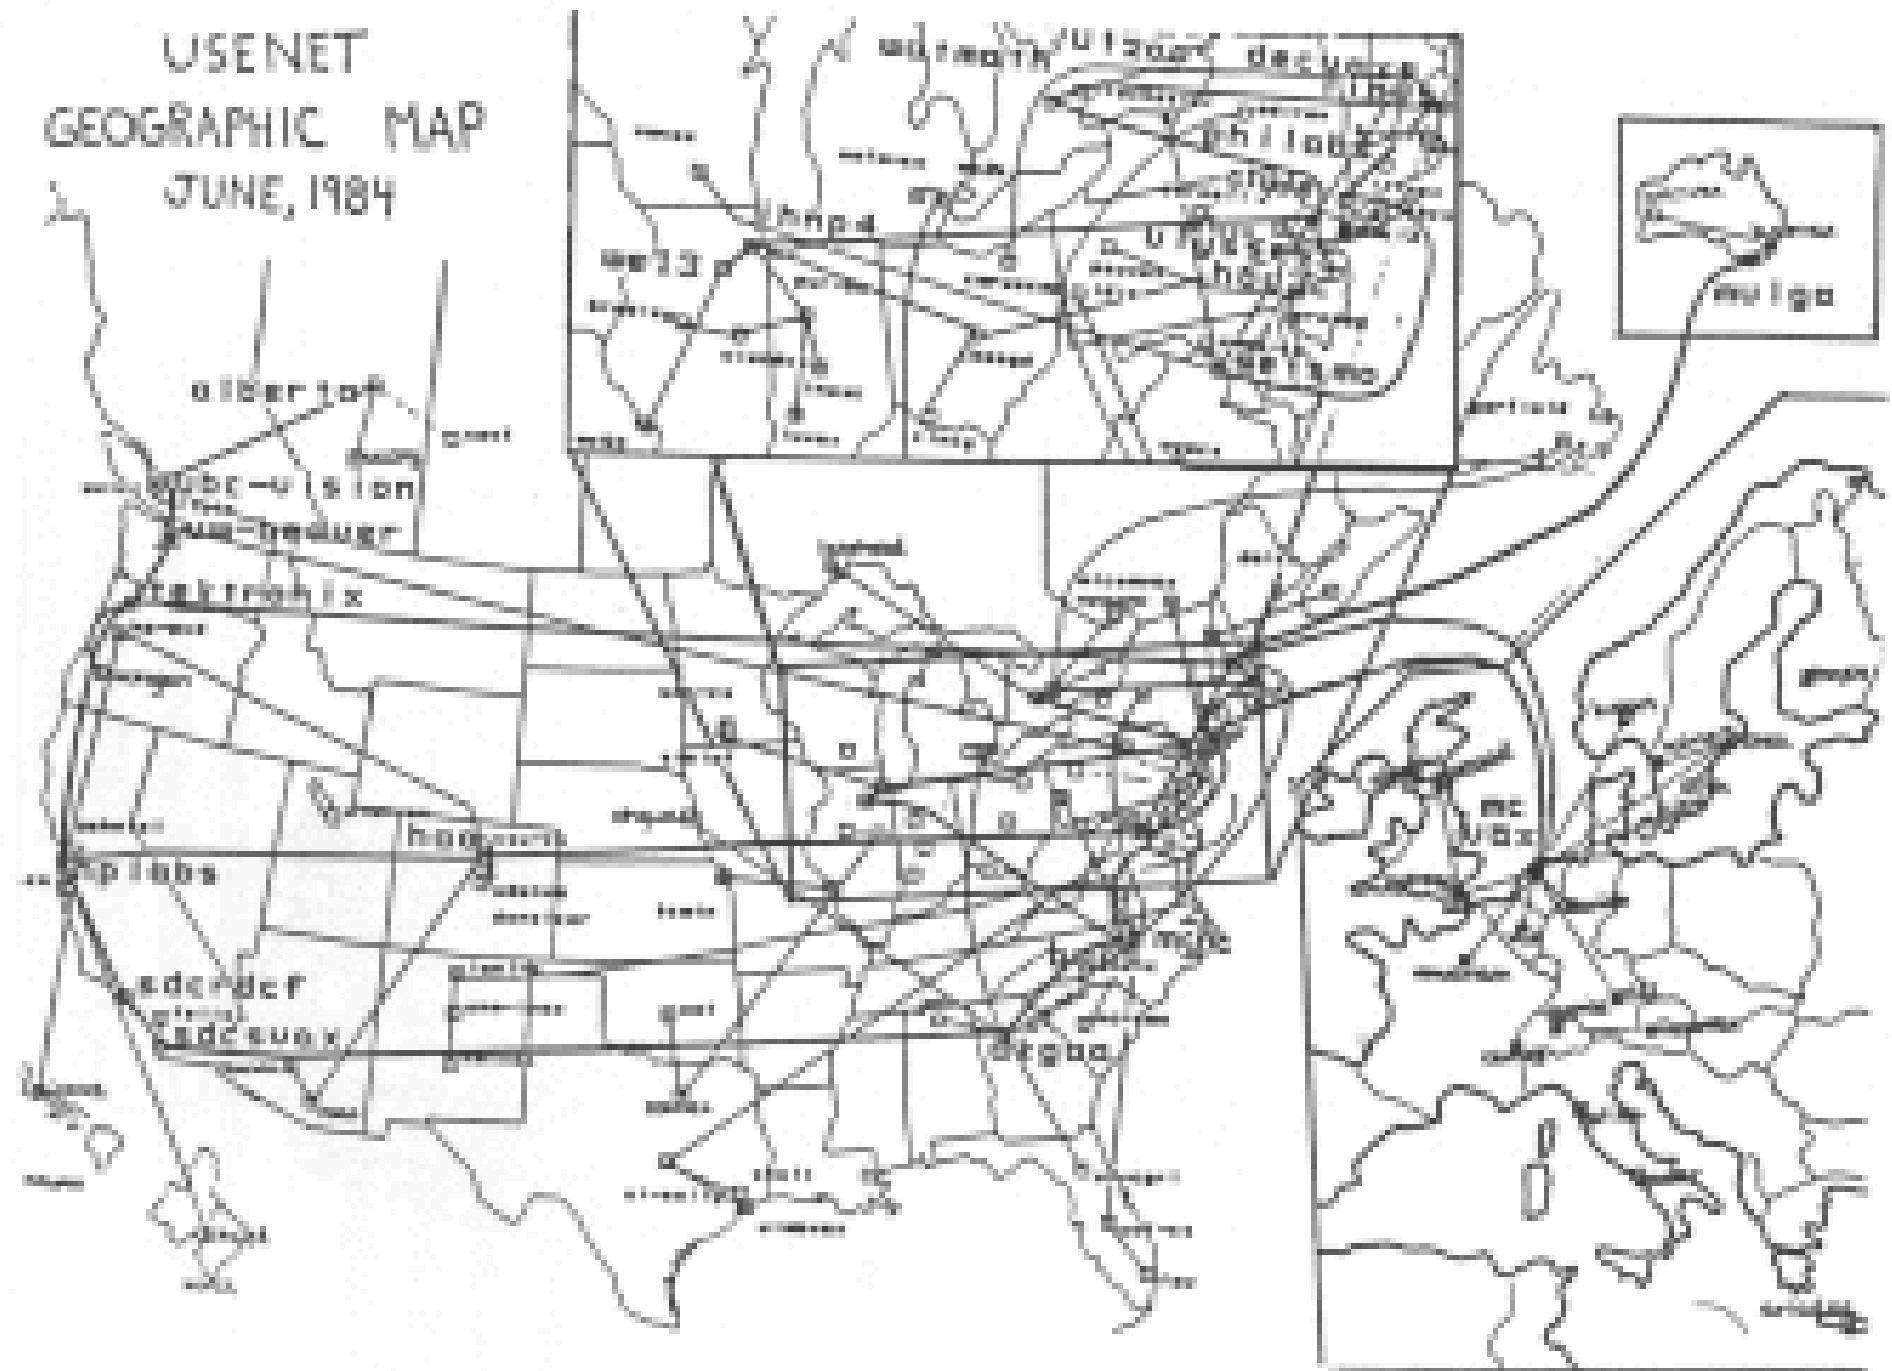
# ARPANET LOGICAL MAP, MARCH 1977



(PLEASE NOTE THAT WHILE THIS MAP SHOWS THE MOST POPULATION OF THE NETWORK ACCORDING TO THE BEST INFORMATION OBTAINABLE, NO CLAIM CAN BE MADE FOR ITS ACCURACY)

NAMES SHOWN ARE IMP NAMES, NOT (NECESSARILY) HOST NAMES


USENET  
GEOGRAPHIC MAP  
JUNE, 1984




# The Internet

Source: LUMETA  
Date: 2014-01-01  
Version: 1.0  
Author: LUMETA  
License: CC BY-NC-SA

- Legend
- ISP group
  - AT&T group
  - Core group
  - Other companies
  - Other backbone





<http://www.opte.org/maps/>


SOME RIGHTS RESERVED






NYC

<http://www.darksky.org/key/ny.html>




World population




## In 2007, 50% of the world is urban

- It was 3% in 1800
- 14% in 1900
- 61% expected in 2030








Earth at Night  
More information available at:  
<http://antwrp.gsfc.nasa.gov/apod/ap001127.html>

Astronomy Picture of the Day  
2000 November 27  
<http://antwrp.gsfc.nasa.gov/apod/astropix.html>




<http://www.opte.org/maps/>


SOME RIGHTS RESERVED




**World Population**    **Internet Users**



## Internet Users Growth





<http://www.opte.org/maps/>

SOME RIGHTS RESERVED






NYC










people ARE the city





**Marshall McLuhan**, well-known for coining the expression "**global village**".





**Marshall McLuhan**, well-known for coining the expression "**global village**".

"All media are extensions of some human faculty- psychic or physical"...

"...the wheel is an extension of the foot the book is an extension of the eye, clothing, an extension of the skin,

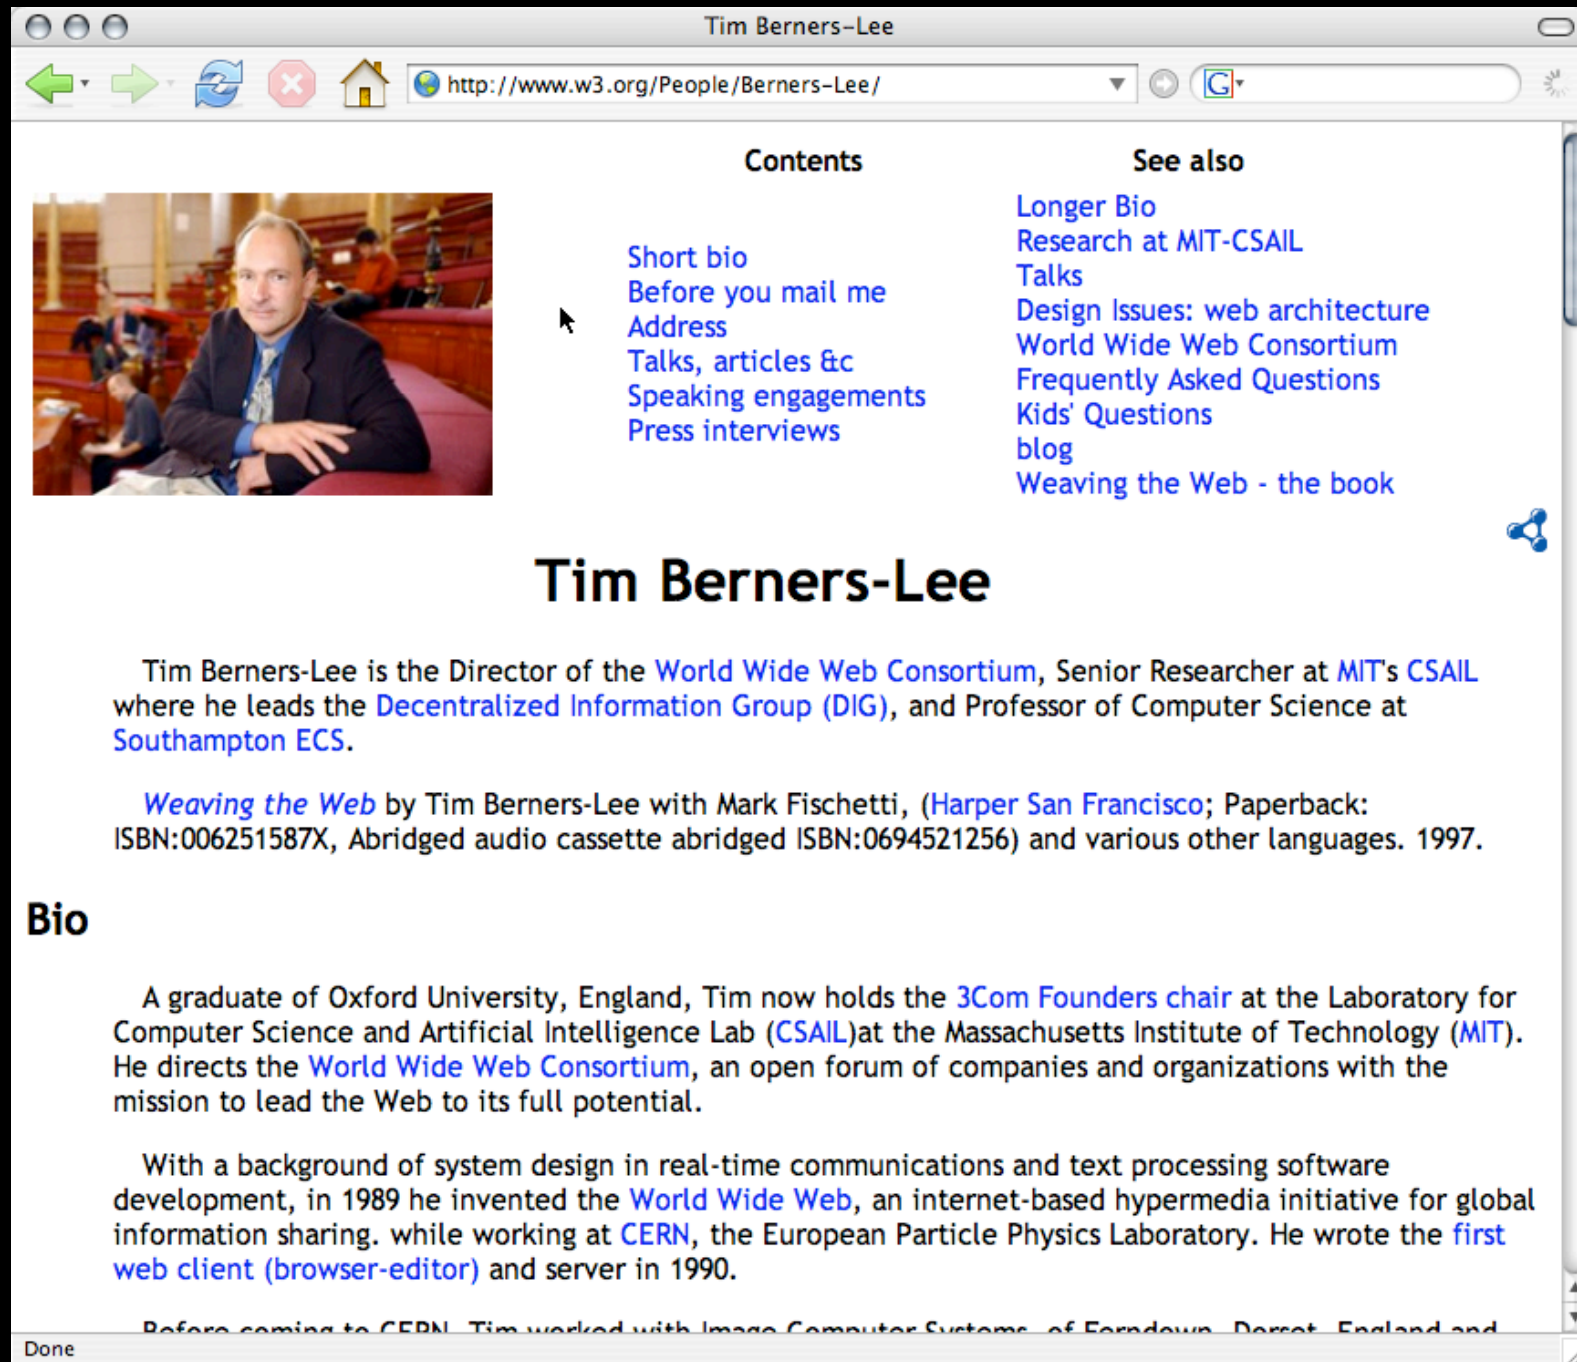
**electric circuitry, an extension of the central nervous system"**

- *The Medium is The Massage*

# people ARE the internet



# a web page:



The screenshot shows a web browser window with the title "Tim Berners-Lee" and the address bar containing "http://www.w3.org/People/Berners-Lee/". The page content includes a navigation menu with "Contents" and "See also" sections. The "Contents" section lists links for "Short bio", "Before you mail me", "Address", "Talks, articles & Speaking engagements", and "Press interviews". The "See also" section lists links for "Longer Bio", "Research at MIT-CSAIL", "Talks", "Design Issues: web architecture", "World Wide Web Consortium", "Frequently Asked Questions", "Kids' Questions", "blog", and "Weaving the Web - the book". A photograph of Tim Berners-Lee is on the left. Below the navigation is a large heading "Tim Berners-Lee" followed by a paragraph describing his roles at MIT and Southampton. A book recommendation for "Weaving the Web" is also present. A "Bio" section follows, detailing his education and work at CERN. The browser's status bar at the bottom shows "Done".


Tim Berners-Lee

Contents

- [Short bio](#)
- [Before you mail me](#)
- [Address](#)
- [Talks, articles & Speaking engagements](#)
- [Press interviews](#)

See also

- [Longer Bio](#)
- [Research at MIT-CSAIL](#)
- [Talks](#)
- [Design Issues: web architecture](#)
- [World Wide Web Consortium](#)
- [Frequently Asked Questions](#)
- [Kids' Questions](#)
- [blog](#)
- [Weaving the Web - the book](#)



## Tim Berners-Lee

Tim Berners-Lee is the Director of the [World Wide Web Consortium](#), Senior Researcher at [MIT's CSAIL](#) where he leads the [Decentralized Information Group \(DIG\)](#), and Professor of Computer Science at [Southampton ECS](#).

*[Weaving the Web](#)* by Tim Berners-Lee with Mark Fischetti, ([Harper San Francisco](#); Paperback: ISBN:006251587X, Abridged audio cassette abridged ISBN:0694521256) and various other languages. 1997.

### Bio

A graduate of Oxford University, England, Tim now holds the [3Com Founders chair](#) at the Laboratory for Computer Science and Artificial Intelligence Lab ([CSAIL](#)) at the Massachusetts Institute of Technology ([MIT](#)). He directs the [World Wide Web Consortium](#), an open forum of companies and organizations with the mission to lead the Web to its full potential.

With a background of system design in real-time communications and text processing software development, in 1989 he invented the [World Wide Web](#), an internet-based hypermedia initiative for global information sharing. while working at [CERN](#), the European Particle Physics Laboratory. He wrote the [first web client \(browser-editor\)](#) and server in 1990.

Before coming to CERN, Tim worked with Image Computer Systems, of Ferndown, Dorset, England and

Done

ebay - New & used electronics, cars, apparel, collectibles, sporting goods & more at low prices

http://www.ebay.com/

WorldClock "The Crack L...ight version AntHdw mob animv VLAN GoodBIN9 searchv goodNIX GoodRSYNC news (12)v lbv asv birdfeedersv >>

home | ebay | site map

Buy Sell My eBay Community Help

Start new search Search

Java™ (Powered by Sun)

Hello, bisquitte! (Sign out)

Whatever it is...you can get it on

WorldClock "The Crack L...ight version AntHdw mob animv VLAN GoodBIN9 searchv goodNIX GoodRSYNC

http://en.wikipedia.org/wiki/Bsd

article discussion edit this page history

Sign in / create account

Specialty Sites

- eBay Express
- eBay Motors
- eBay Stores
- eBay Business
- Hall.com
- Apartments on Rent.com

ending

WIKIPEDIA The Free Encyclopedia

navigation

- Main page
- Community portal
- Featured content
- Current events
- Recent changes
- Random article
- About Wikipedia
- Contact us
- Make a donation
- Help

Berkeley Software Distribution

From Wikipedia, the free encyclopedia

(Redirected from Bsd)


"BSD" redirects here. For other uses, see BSD (disambiguation).

**Berkeley Software Distribution** (**BSD**, sometimes called **Berkeley Unix**) is the Unix derivative distributed by the **University of California, Berkeley**, starting in the **1970s**. The name is also used collectively for the modern descendants of these distributions.

BSD is one of several branches of Unix operating systems. Another one is evolved from **UNIX System V** developed by AT&T's Unix System Development Labs. A third consists of the **GNU/Linux** operating systems which draw from Unix System V and BSD, as well as **Plan9**, and non-UNIX operating systems.

BSD was widely identified with the versions of Unix available for workstation-class systems. This can be attributed to the ease with which it could be licensed and the familiarity it found among the founders of many technology companies during the 1980s. This familiarity often came from using similar systems—notably DEC's **Ultron** and Sun's **SunOS**—during their education. While BSD itself was largely superseded by the **System V Release 4** and **OSF/1** systems in the 1990s (both of which incorporated BSD code), in recent years modified **open source** versions of the codebase (mostly derived from 4.4BSD-Lite) have seen increasing use and development.

BSD Unix



Website: N/A

Company/developer: CSRG, UC Berkeley

OS family: Unix

Source model: Open source

Latest stable release: 4.4-Lite2 / 1995

Kernel type: Monolithic

License: BSD license

Contents [hide]

- 1.1 PDP-11 beginnings
- 1.2 VAX versions
- 1.3 4.3BSD
- 1.4 Net/2 and legal troubles
- 1.5 4.4BSD and descendants

WorldClock "The Crack L...ight version AntHdw mob animv VLAN GoodBIN9 searchv goodNIX GoodRSYNC news (12)v lbv asv birdfeedersv >>

http://www.hopstop.com/

WorldClock "The Crack L...ight version AntHdw mob animv VLAN GoodBIN9 searchv goodNIX GoodRSYNC news (12)v lbv asv birdfeedersv >>

HOPSTOP YOUR CITY TRANSIT GUIDE

Online Subway and Bus Directions

The Paul McGhee Division of NYU. A program devoted to adult students.

CLICK HERE TO LEARN MORE

Home Directions Maps City Guide Trips Ratings Mobile About Us My HopStop

Directions: New York

Suggested Route

Starting & Destination Addresses

- 304 PARK AVE S, Manhattan
- 420 9TH AVE, Manhattan

Options

Transportation Mode: Subway only

Walking/Transfer Preference: Less street walking/More transfers

Route Plan

Route	Time
Start out going North on Park Ave S towards E 23rd St	0.2
Entrance near intersection of E 23rd St and Park Ave S	2.0
Take the 6 train from 23 Street station heading Uptown / Pelham Bay Park	5.0
Pass 28 Street	1.4
Pass 33 Street	1.4

www.myspace.com/laserbeast

CONGRATULATIONS! You've Been Selected! FREE \$2500

Click Here to Claim!

MySpace

Home | Browse | Search | Invite | Film | Mail | Blog | Favorites | Forum | Groups | Events

MYSPACE MUSIC

Lightning Bolt

Drum & Bass / Other / Experimental

Providence, Rhode Island United States

Profile Views: 152113

Last Login: 3/7/2007

View My: Pics | Videos

Contacting Lightning Bolt

- Send Message
- Add to Friends
- Instant Message
- Add to Group
- Forward to Friend
- Add to Favorites
- Block User
- Rank User

MySpace URL: http://www.myspace.com/laserbeast

Lightning Bolt: General Info

- Member Since: 12/20/2005
- Band Website: laserbeast.com/

Gmail - [nycbug-talk] AsiaBSDCon, March 8-11

http://mail.google.com/mail/?auth=DQAAAGwAAChRW...

Google Gmail Calendar Photos Docs & Spreadsheets Groups all my services >

dot.ike@gmail.com | Settings | Help | Sign out

nycbug

Search Mail Search the Web Show search options Create a filter

Harlem Success Academy - www.harlemsuccess.org - Join Education Reform Movement Apply N

[nycbug-talk] AsiaBSDCon, March 8-11 bsdlists

Isaac Levy to NYCBUG Feb 12

Hey All,

The fine folks running AsiaBSDCon are looking to spread the word about the conference, the schedule is now online! Pretty excellent lineup, their tutorial sessions all look pretty hardcore. It seems to me the heavy focus of the conference is of course on BSD's strength in ISP and heavy network contexts- (no KDE tuning sessions? What?!).

Also, of note, lots of IPv6 stuff at this conference. (in \*the\*

New window Print all Turn off highlighting

Sponsored Links (feedback)

Tokyo Flights from \$423 Discounted Fares to Japan. Search, Book and Save Big! www.BT-STORE.com

Cheap Japan Tickets

'web pages', or are they...

# software applications?

zopelists (18158) [Edit labels](#)

▼ Invite a friend  
Give Gmail to:  
  
Send Invite 97 left  
[preview invite](#)

% NYC\*BUG talk mailing list  
<http://lists.nycbug.org/mailman/listinfo/talk>  
%Be sure to check out our Jobs and NYCBUG-announce lists  
%We meet the first Wednesday of the month

Reply Reply to all Forward Invite Isaac to Gmail

Send Save Now Discard

To: Isaac Levy <ike@lesmuug.org>

[Add Cc](#) | [Add Bcc](#) | [Edit Subject](#) | [Attach a file](#)

**B** *I* U  $\mathcal{F}$   $\tau$   $\text{T}$   $\text{T}$  [Check spelling](#) ▼

[Plain text](#)

Are my expectations of web applications **growing**?

On 2/12/07, **Isaac Levy** <ike@lesmuug.org> wrote:  
Hey All,

The fine folks running AsiaBSDCon are looking to spread the word about the conference, the schedule is now online! Pretty excellent lineup, their tutorial sessions all look pretty hardcore. It seems to me the heavy focus of the conference is of course on BSD's strength in ISP and heavy network contexts- (no KDE tuning sessions? What?!).

Also, of note, lots of IPV6 stuff at this conference, (in \*the\* country which actually runs \*production\* grade IPV6 infrastructure!)

Send Save Now Discard

[Back to Search results](#) Report Spam Delete More actions... [Newer 30 of thousands Older](#)





homepages?  
(raising user expectations of the internet)

www.myspace.com/laserbeast

http://www.myspace.com/laserbeast


**CONGRATULATIONS!**  
**You've Been Selected! FREE 42" Plasma TV**  
[Click Here to Claim!](#) \*details apply

MySpace Search powered by Google

Home | Browse | Search | Invite | Film | Mail | Blog | Favorites | Forum | Groups | Events | Videos | Music | Comedy | Classifieds

**MYSPEACE MUSIC** Music Videos | Directory | Search | Top Artists | Shows | Music Forums | Music Classifieds | Artist Signup

**Lightning Bolt**  
Drum & Bass / Other / Experimental



Providence, Rhode Island  
United States

Profile Views: 152113

Last Login: 3/7/2007

View My: [Pics](#) | [Videos](#)

**Contacting Lightning Bolt**

Send Message	Forward to Friend
Add to Friends	Add to Favorites
Instant Message	Block User
Add to Group	Rank User

**MySpace URL:**  
<http://www.myspace.com/laserbeast>

**Lightning Bolt: General Info**

Member Since	12/20/2005
Band Website	<a href="http://laserbeast.com/">laserbeast.com/</a>

**Captain Caveman**  
Lightning Bolt  
playing

00:04

**Total Plays: 355927**    **Downloads Today: 0**    **Plays Today: 1370**

[Dead Cowboy](#) Plays: 24542  
Download | Rate | Comments | Lyrics | Add

[Dracula Mountain](#) Plays: 25570  
Download | Rate | Comments | Lyrics | Add

[Rotator](#) Plays: 21896  
Download | Rate | Comments | Lyrics | Add

[Captain Caveman](#) Plays: 21110  
Download | Rate | Comments | Lyrics | Add

Hypermagic Mountain  
2005 Load

**STANDALONE PLAYER**

**Lightning Bolt's Latest Blog Entry** [[Subscribe to this Blog](#)]

tours ([view more](#))

Fan Art!! ([view more](#))

We'll Be Back! ([view more](#))

[[View All Blog Entries](#)]

**About Lightning Bolt**

**Lightning Bolt's Friend Space**

Transferring data from cache03-music02.myspacecdn.com...



**Rocinha**



*Stewart Brand, Currently working with the Long Now Foundation*

how did iMeme happen?

# Typical ISP Setup (virtual hosting)

1999:  
bleak options  
(unless you  
want LAMP)

Affordable, full service web hosting packages.

Start a domain search:  com  24/7 Sales & Support: (480)505-8877 Hot Deals! SALE!

Go Daddy .COM Make a .com name with us!®

SPECIAL OFFER for RegisterFly.com Customers!

SEE OUR SUPER BOWL TV AD & WEB-ONLY VERSION!

DISCOUNT DOMAIN CLUB Exclusive Discounts on domain purchases!

BobParsons.com The Registerly scandal. Missing funds, escort services, liposuction & a Miami penthouse.

Domains | Hosting & Servers | Email | Site Builders | Business | SSL Certificates | Domain Auctions | Reseller Plans

Go Daddy Home > Hosting > Checkout

## Hosting Plans

The world's largest hostname provider!

**FREE Setup!**  
Now up to 200 GB storage & 2000 GB transfer!  
99.9% Uptime!\*

GoDaddy.com, the World's largest hostname provider\*, offers plans ideal for most individuals and small businesses. We're the affordable, reliable place to host your site – with no ad banners or pop-ups, and 24/7 support and [free access to our Metropolis Hosting Community](#). And unlike the competition, there's no set up fee and no annual commitment required.

**EXCLUSIVE metropolis hosting community**  
Over 30 FREE applications help you get the most out of your hosting!

My Account Logout  
Customer # or Login name:  
Password:  
[Forgot your password?](#)  
[Assign an AccountExec!](#)  
Secure Login

My Renewals  
My Cart Currency:   
My Email  
My News  
My Radio News  
Earn Money - Be A Reseller! \$

### Choose your hosting plan!

See complete [Hosting Plan comparison chart](#)

Linux  Select Windows or Linux:

#### Economy Plan:

- 5 GB Space • 250 GB Transfer
- 500 Email Accounts • **FREE!** Software
- 10 MySQL Databases • 50 Email Forwards
- Forums, Blogging, Photo Galleries • No ads
- **EXCLUSIVE!** [Metropolis Hosting Community](#)

2 months: **Just \$3.99/mo**  
12 months: **Just \$3.59/mo SAVE 10%!**  
24 months: **Just \$3.19/mo SAVE 20%!**

#### Deluxe Plan:

- 100 GB Space • 1,000 GB Transfer
- 1,000 Email Accounts • **FREE!** Software
- 25 MySQL Databases • Unlimited Email Forwards
- Forums, Blogging, Photo Galleries • No ads
- **EXCLUSIVE!** [Metropolis Hosting Community](#)
- **FREE!** [\\$25 Google® AdWords® Credit!](#)

1 month: **Just \$6.99/mo**  
12 months: **Just \$6.29/mo SAVE 10%!**  
24 months: **Just \$5.59/mo SAVE 20%!**

#### Premium Plan:

- 200 GB Space • 2,000 GB Transfer
- 2,000 Email Accounts • **FREE!** Software
- 50 MySQL Databases • Unlimited Email Forwards
- Forums, Blogging, Photo Galleries • No ads
- **EXCLUSIVE!** [Metropolis Hosting Community](#)
- **FREE!** [SSL Certificate, \\$19.99 value!](#)
- **FREE!** [\\$25 Google® AdWords® Credit!](#)

1 month: **Just \$14.99/mo**  
12 months: **Just \$13.49/mo SAVE 10%!**  
24 months: **Just \$11.99/mo SAVE 20%!**

Done [www.godaddy.com](#)

# Typical ISP Setup (virtual hosting)



```
../vhosts  
  /customer1  
    /www  
    /cgi_bin  
  /customer2  
    /www  
    /cgi_bin  
  /customer3  
    /www  
    /cgi_bin
```



```
#customer1  
#customer2  
#customer3  
#customer4  
#customer5  
#customer6  
#customer7
```



# The Utopian City?



“A House Is  
A Machine  
For Living In”  
Le Corbusier



**Jane Jacobs,**  
well-known for working to  
maintain diverse and  
meaningful NYC urban life.

With regard to the ideals of Le Corbusier's ideas for city design:

**“...As in all Utopias, the right to have plans of any significance belonged only to the planners in charge.”**

*- The Death and Life of Great American Cities*

# many shapes and scales...

(growth as complex land use, not merely size)



apartment building

(note mixed use retail)



apartment building

(misc west village)




skyscraper

(Empire State Building)



# UNIX



# Dedicated Hosting?



**1999:**

Dedicated hosting can be prohibitively **expensive**, (and often, a waste of hardware resources!)





## **Doug McIlroy,**

the inventor of Unix pipes, part of early UNIX development.


“This is the Unix philosophy:

**Write programs that do one thing and do it well.**

Write programs to work together. Write programs to handle text streams, because that is a universal interface.”

- Peter H. Salus. *A Quarter-Century of Unix*. Addison-Wesley. 1994. ISBN 0-201-54777-5


# The Otis Elevator 1864 (it does one thing)




Elijah Otis - The Elevator

“In the era of the staircase all floors above the second were considered unfit for commercial purposes, and all those above the fifth, uninhabitable.”

“Since the 1870s in Manhattan, the elevator has been the great emancipator of all horizontal surfaces above the ground floor.”








Poul-Henning Kamp - jail(8)




Robert Watson too...



Elijah Otis - The Elevator






Poul-Henning Kamp - jail(8)

jail(8)  
1998

**(it does one thing)**

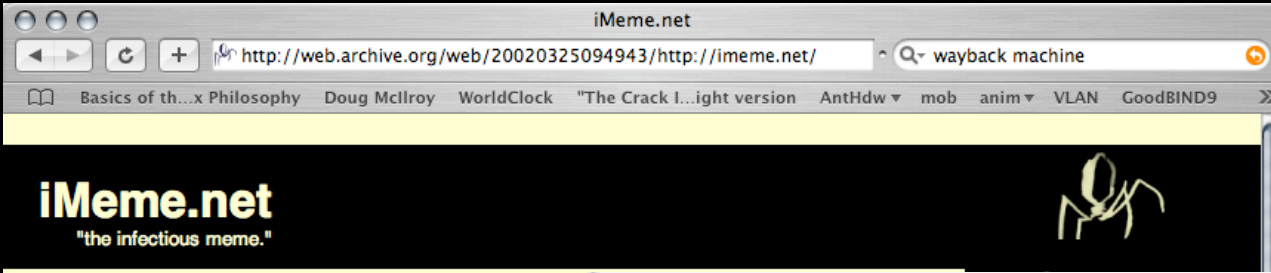


Robert Watson worked on jail(8) as well...



jailing server

A screenshot of a web browser displaying the imeme.net website. The browser's address bar shows the URL: http://web.archive.org/web/20000816090203/http://imeme.net/. The page title is "Web Hosting". The main content area features a light blue background with the text: "Welcome to [imeme.net](http://imeme.net), an Internet Application Service Provider." Below this, a dark box contains the text "Everything through the web". Further down, it says "This is what you will get:" followed by the following links and offers: [you.imeme.net](http://you.imeme.net), [you@imeme.net](mailto:you@imeme.net), [friends@lists.imeme.net](mailto:friends@lists.imeme.net), and "100 MB Space". A dark box at the bottom of the main content area states "Only \$10 a month No Setup Fee". On the right side, there is a sidebar with a dark background and white text. It is titled "Technologies" and lists three categories: "Zope" (with a bullet point: "Dynamic Web Presence", "The feature for the rest of us", "Work in Versions", "Included Products"), "WorldPilot" (with a bullet point: "IMAP Email account", "Personal Information Manager"), and "Mailman" (with a bullet point: "Personal Mailing List"). The browser's status bar at the bottom left shows the word "Done".



**iMeme.net**  
"the infectious meme."



## Company Information

### Who We Are

iMeme is a Zope hosting provider dedicated to Open Source and conveniences for Zope developers. We specialize in hosted Zope developers community. Each iMeme account is overseen directly by an owner of the company.

### News

Helium is back up after approximately 1h of downtime due to a page fault; we are investigating to attempt to determine the cause.

[\[More ...\]](#)

### What We Offer

Our standard webhosting package provides you with the following features:

#### Individual Zope instance

Control every aspect of your Zope. You can use experimental versions of products, build and install your own products with the latest stable version of Zope by default. You can also provide future scalability and enhanced access to your data.

#### Root Access

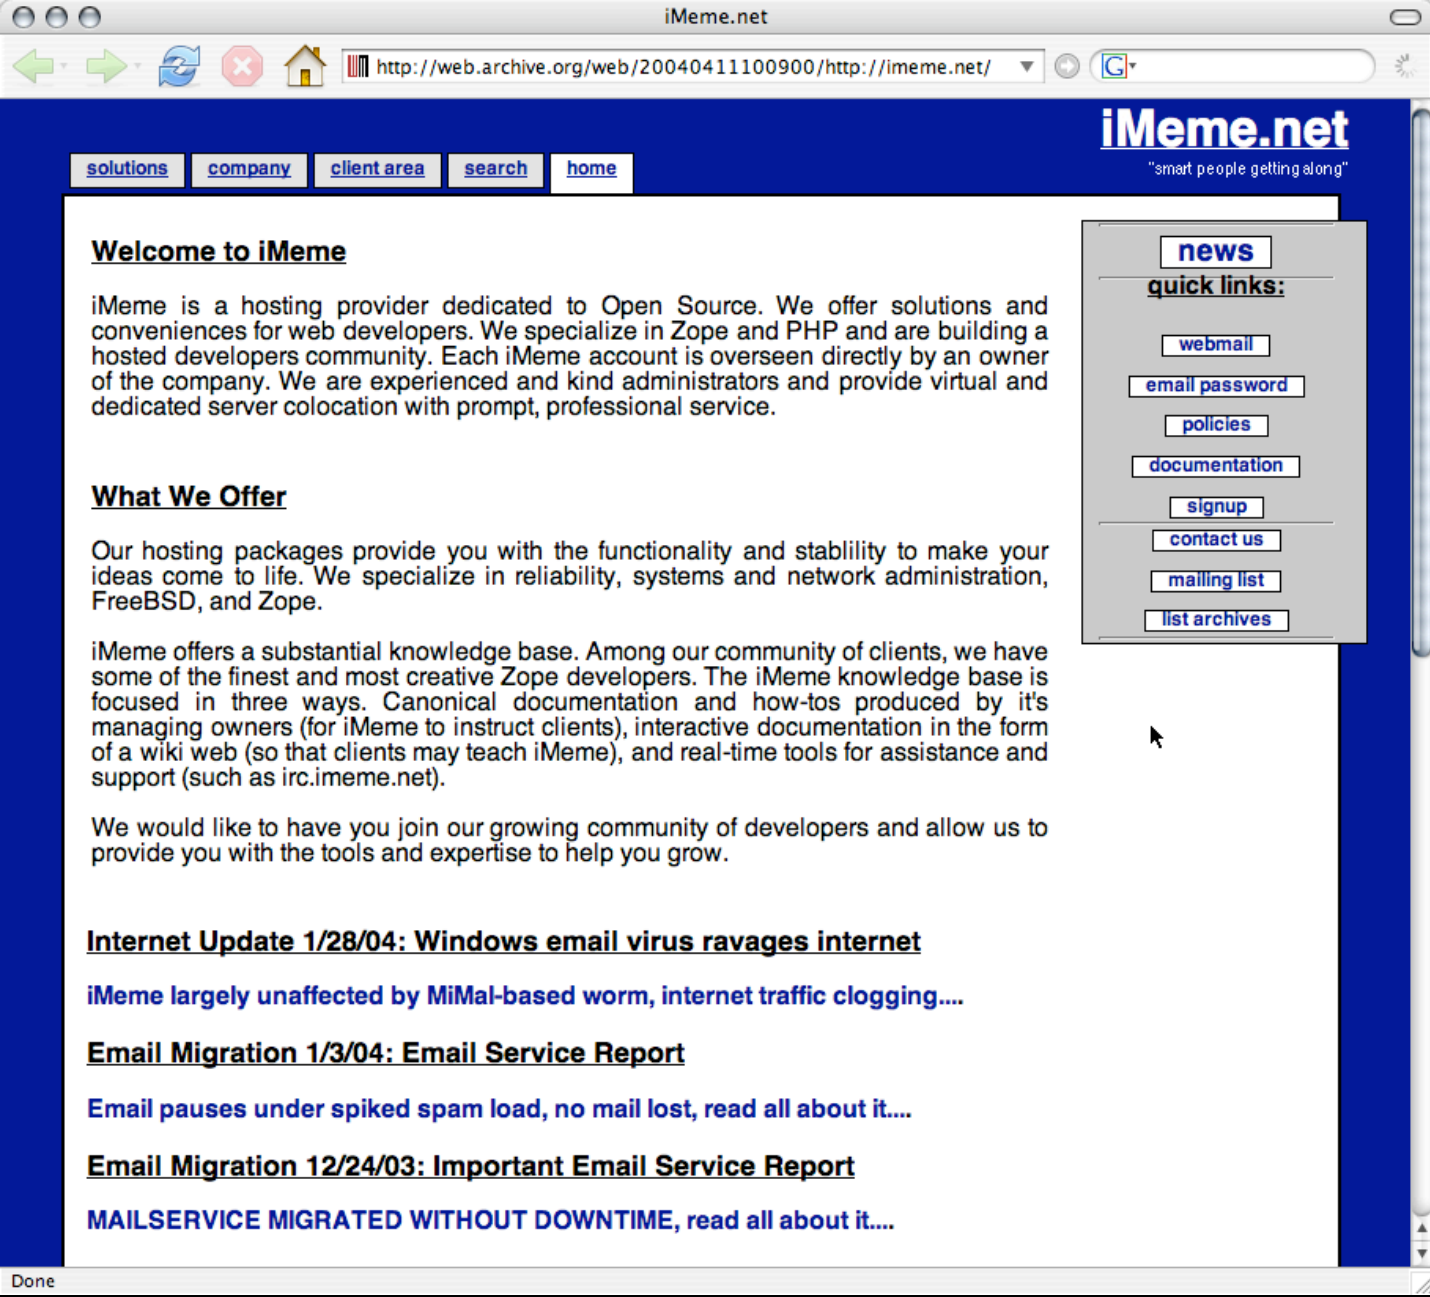
Root access within your account. This gives you the ability to manage permissions, and have a full development environment.

#### Filesystem access via SSH

Securely access the server running your web presence and develop custom applications.

#### Unlimited IMAP email addresses

Provide your client with email addresses at their domain. Email for retention of read status even with multiple accounts amongst multiple parties.



**iMeme.net**  
"smart people getting along"

- [solutions](#)
- [company](#)
- [client area](#)
- [search](#)
- [home](#)

## Welcome to iMeme

iMeme is a hosting provider dedicated to Open Source. We offer solutions and conveniences for web developers. We specialize in Zope and PHP and are building a hosted developers community. Each iMeme account is overseen directly by an owner of the company. We are experienced and kind administrators and provide virtual and dedicated server colocation with prompt, professional service.

### What We Offer

Our hosting packages provide you with the functionality and stability to make your ideas come to life. We specialize in reliability, systems and network administration, FreeBSD, and Zope.

iMeme offers a substantial knowledge base. Among our community of clients, we have some of the finest and most creative Zope developers. The iMeme knowledge base is focused in three ways. Canonical documentation and how-tos produced by it's managing owners (for iMeme to instruct clients), interactive documentation in the form of a wiki web (so that clients may teach iMeme), and real-time tools for assistance and support (such as irc.imeme.net).

We would like to have you join our growing community of developers and allow us to provide you with the tools and expertise to help you grow.

### Internet Update 1/28/04: Windows email virus ravages internet

[iMeme largely unaffected by MiMal-based worm, internet traffic clogging...](#)

### Email Migration 1/3/04: Email Service Report

[Email pauses under spiked spam load, no mail lost, read all about it...](#)

### Email Migration 12/24/03: Important Email Service Report

[MAILSERVICE MIGRATED WITHOUT DOWNTIME, read all about it...](#)

**news**

**quick links:**

- [webmail](#)
- [email password](#)
- [policies](#)
- [documentation](#)
- [signup](#)
- [contact us](#)
- [mailing list](#)
- [list archives](#)

There were 2 errors opening the page. For more information, choose Activity from the menu.



# The real world of iMeme users:

- hacker: “I want to compile LISP”
- undergraduate sociology student: “I want to install ‘Foo’ blog software, it’s PHP and the instructions say I need to run Cron”
- web designer: “I want to run an http server on port 8080”
- business owner: “I want to run Foo web application for my business.”
- A community leader: “I want to run Mailman List Manager”
- 13 year old hacker: “I want to run both an IRC and jabber server for my friends”.

**Most iMeme users simply wanted to use Python/Zope.**



# The UNIX Time-Sharing System\*

*D. M. Ritchie and K. Thompson*


## ABSTRACT


Unix is a general-purpose, multi-user, interactive operating system for the larger Digital Equipment Corporation PDP-11 and the Interdata 8/32 computers. It offers a number of features seldom found even in larger operating systems, including

- i A hierarchical file system incorporating demountable volumes,
- ii Compatible file, device, and inter-process I/O,
- iii The ability to initiate asynchronous processes,
- iv System command language selectable on a per-user basis,
- v Over 100 subsystems including a dozen languages,
- vi High degree of portability.


This paper discusses the nature and implementation of the file system and of the user command interface.

**NOTE:** \* Copyright 1974, Association for Computing Machinery, Inc., reprinted by permission. This electronic edition of this paper is a reprint of the version appearing in *The Bell System Technical Journal* 57 no. 6, part 2 (July-August 1978). In turn, that was a revised version of an article that appeared in *Communications of the ACM* 17, No. 7 (July 1974), pp. 365-375. That article was a revised version of a







jailing server




apartment building  
(misc west village)




jailing server




skyscraper  
(Empire State Building)





# Internet Service Provider (ISP)



**Contemporary Datacenter**



# Technical Application



# management issues

- lost jail?
  - [hostname lockdown]
- resource attacks
  - disks full
    - [partitions, disk images]
  - fork bombs, memory hogs
    - [securelevels, login.conf]
  - process control
- direct driver access
  - [flags to mount devfs, procfs]

# management issues

- lost jail?
  - [hostname lockdown]
- resource attacks
  - disks full
    - [partitions, disk images]
  - fork bombs, memory hogs
    - [securelevels, login.conf]
  - process control
- direct driver access
  - [flags to mount devfs, procfs]



# General Specs

- High Density 2u (and then 1u) servers
- we aimed to run 50 jails per box
- provisioned 4gb of disk space
- 100mb of what we called 'process space', the amalgamation of memory and cpu usage
- Bandwidth was rarely an issue, very basic QOS oriented throttling

# General Specs

- Company Scale in 2003
  - 1,000 domains
  - 480 jails
- less than half a rack of servers (24u)
- 3 owners, (2 of us ran operations)

# Managing Disk Use

- we ran scripts from the host server which simply used `du`, and shoved the output into MySQL databases
- we then automated the process of enforcing policies of charging for extra disk usage. (simple cron jobs to email users, change their bill, etc...)
- iMeme gave 1 month of 'grace time' to trim disk usage, sometimes logfiles would explode, or users would accidentally consume undue disk space- and we felt this was a simple buffer our people appreciated

# Managing Disk Use

- FreeBSD file-backed memory disks (disk images)
  - `mdconfig(8)`, consult FreeBSD handbook
- Extreme I/O penalty, which is the bottleneck for jailing already!
- disk images consume RAM!
- encrypted disk images could be used as well...

# Managing Memory

- Memory and CPU usage polled on a regular basis for each jail, we called it 'Process Space'
- Shell scripts were originally setup to run as cron jobs inside each jail, which took cumulative memory consumption and cpu usage by parsing ps(1) output inside a given jail
- While iMeme originally ran these scripts inside of each jailed system, outputting totals to text files in /jail/dir/var/log/,
  - this carried the risk that a user could (trivially) bypass this system to avoid increased billing or otherwise.



# Managing Memory

- iMeme moved this system out of the jails themselves, to the host system with new jailing features in FreeBSD 5.x- (6.x)
  - one can list/kill processes based on the jail id, information available to ps, and processes listed in the /proc filesystem.
  - jail(8) info is noted in many common utilities 5.x onward

# Managing Processes

- process restrictions were then handled neatly using `renice(8)`
- shell scripts polled every 5 minutes for ‘hog’ processes, which were logged to SQL
  - if the process wouldn’t behave in 5 minutes, it was `reniced`
- crude but wildly successful

# Managing Processes

- Fork bombs were still a threat, but rare.
- Fix not feasible, required booting host server with escalated securelevel

# Managing Processes

- from FreeBSD 5.x onward, easy to fix:
  - each jail can be set to start with an escalated securelevel
  - maxprocs could be locked for a jail
  - chflags(2) disabled in jails via host sysctl settings

# Network Mention

Far outside the scope of this material, however, it is worth mentioning one thing:

at iMeme, **each jailing host** server was conceptually treated like a **network border** or gateway, with routing and filtering tasks carried out inside the machine.

# Network Mention

With that, we ran NAT for our external IP blocks, and mapped addresses to our jails- which all ran using a private netblock, (192.168.x.x).

# Network Mention

- ipfw and dummynet for simple fair-share traffic shaping
- ipfw performed strongly for putting out fires
- Current Diversiform systems, (me), pf replaces ipfw all around- though high-volume scale usage has not come close to 'iMeme levels'

# Hostname Annoyance

- 4.x jailing relied heavily on a jailed hostname for userland jail identification
  - jailed user could change hostname, making it VERY hard to find
- FreeBSD 5.x solved this problem by:
  - pinning a 'jail id' to each process on the system
  - providing a sysctl feature to lock down the ability to change hostnames within a jail



# deploying new systems

- iMeme used CVS to manage jail 'skeletons'
- Tarball deployment is your friend.  
clean, simple, reliable.
  - be aware of dev/proc mounts
  - be aware of symlinks
  - be aware of permissions, (and file flags)
- use FreeBSD Ports Mechanism?!  
(**not** for the ports collection proper, that's  
insanely presumptuous, [borderline insane])

# upgrading jailed systems

- Simply use buildworld, (FROM HOST SYSTEM),
- give buildworld the DESTDIR flag, with a jail's userland path
- just follow the handbook: [http://www.freebsd.org/doc/en\\_US.ISO8859-1/books/handbook/makeworld.html](http://www.freebsd.org/doc/en_US.ISO8859-1/books/handbook/makeworld.html)

# managing jailed systems

- iMeme website, where users bought jailed systems, and managed their account and billing, was all written in Zope, and had PHP elements added over time.  
**This could have been any web technology.**
- iMeme kept a 'Master Record Server' (redundant), running a MySQL database with all jail/customer data
  - jailed system details
  - resource usage
  - billing and contact information

# jumping jails, backups

- every server mounted every other server
- dirty NFS setup (reliable, but clunky)
  - single point of failure for cluster of host servers
- If a jailing host server died:
  - it's jailed systems could then be rapidly re-distributed across the entire cluster
  - required intelligent Administrator intervention

# best practices

- ssh into jails to manage their processes!  
[ jexec(8) jail\_attach(2) have consequences?!]
- You always can see the jailed filesystem/  
userland from host server, be careful.
- Design jailing host system carefully, be  
creative with core UNIX utilities.
- Use your highest secure practices for host  
server...

# best practices

(stay away from rc jail scripts in ISP context!!!)

```
#!/bin/sh
```

```
# simple, complete script to start a jail.
```

```
# define the absolute path to the jail,  
J=/usr/local/jails/jailed.userland.directory
```

```
# define the ip address for the jail,  
I=10.0.1.192
```

```
# define a hostname,  
H=fqdn.com
```

```
ifconfig en0 inet alias $I/32
```

```
mount -t procfs proc $J/proc  
mount_devfs devfs $J/dev
```


```
## add additional flags to mount_devfs, to hide unnecessary devices!!!  
## check the man page for mount_devfs
```

```
jail $J $H $I /bin/sh /etc/rc
```

# important utilities

- 4.x FreeBSD, jps, jkill, jtop (ports)
- 5.x, 6.x, onward builtin ps, kill
  - !plus jls(8), jexec(8) jattach(2), sysctl features for jailing
- Design your jailing system carefully, be creative (note about nullfs, devfs)
- additionally, handy: pstree, xtail, disk images via mdconfig


## Internet Users Growth






# future directions (for me)...

and if I did it again,  
(a full blown ISP):



- GEOM, GGated over NFS?
- CARP, from PF/OpenBSD - makes life easier
- more NAS/SAN support (GEOM, ggated)
- systrace?! <http://techie.devnull.cz/systrace/>
- Experiment with SysJail, (OpenBSD, NetBSD!)

# current directions (for me)...



## **CATCHING MY BREATH AFTER iMeme!!!**

- working freelance from NYC
- running highly redundant jailed systems for clients
- always hacking around with jail stuff...

# **NMCBUG**

NEW YORK CITY \*BSD USER GROUP


**Jailing Party - jail(8)  
Friday-Monday, Sept. 3-6 2004  
Next time we get more resilient hardware...  
Thanks everyone for all the fun!**



# Special Thanks:

**NMCBUG**

NEW YORK CITY \*BSD USER GROUP



**diversafirm inc.**



# Special Thanks:



**wintermute** (partner at iMeme), taught me to jail(8).




**Poul-Henning Kamp** wrote the jail feature for R&D Associates <http://www.rndassociates.com/> who contributed it to FreeBSD around 1998.



**Robert Watson** wrote the extended documentation, found a few bugs, added a few new features, and cleaned up the userland jail environment.



**Q&A?**



**jail(8)**