# LLVM on the Web

## Using Portable Native Client to run Clang/LLVM in the Browser

February 1, 2015

Brad Nelson
**@flagxor**

# Motivation

Author Ron Reuter (aka MainByte) I created the photo/file and it is released to the public domain.

# Motivation

- I hate configuring new computers
- I don’t trust all the programs I run
- I’m a programmer

# Goals

- Develop in the Browser
- Make it SECURE
- Make it FAST

# What I Need

# What I Need

## Native Code

# Native Code on the Web?

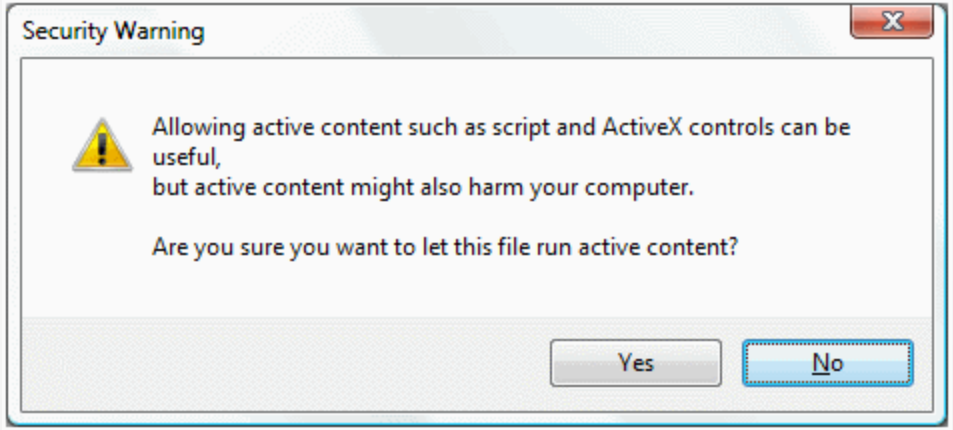

# Native Client

- Native code as secure as JavaScript
- Verifiably safe
- Portable

# Jargon

**Na**tive **Cl**ient → **NaCl**

**P**ortable **Na**tive **Cl**ient → **PNaCl**

**P**epper **P**lugin **API** → **PPAPI**

# Native Client

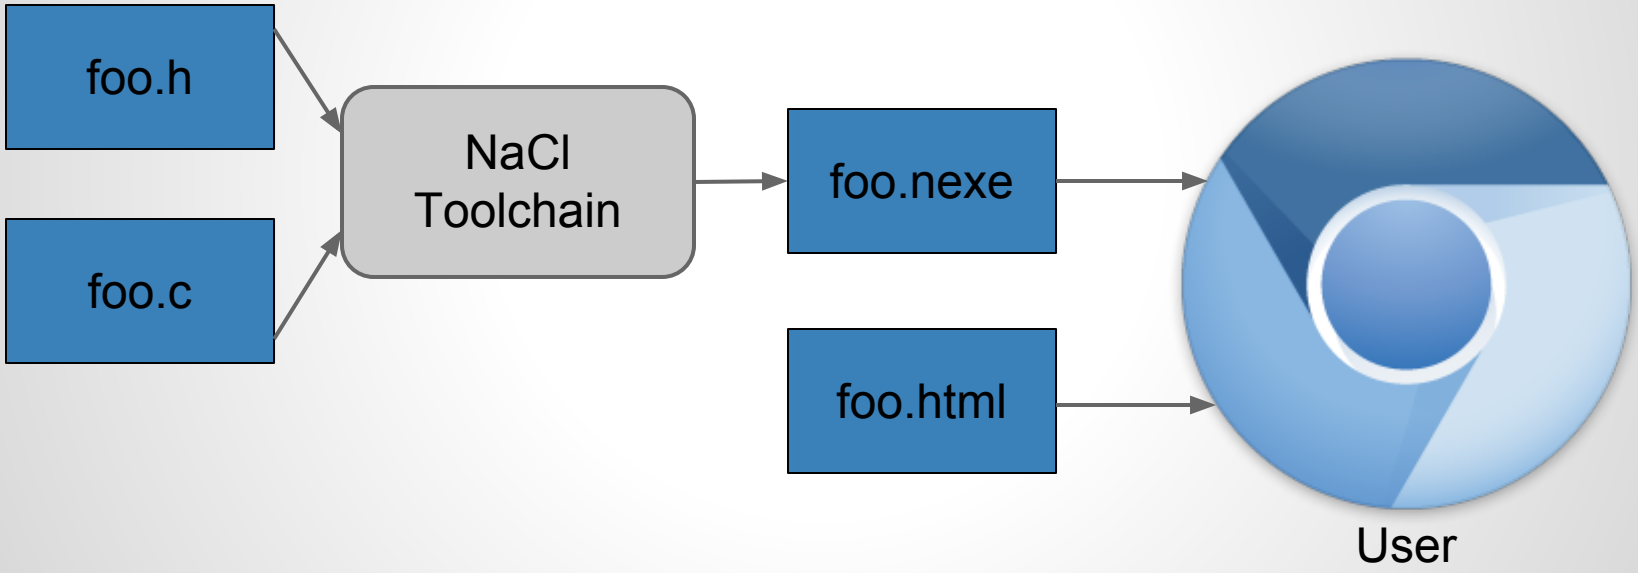

# Two Sandboxes!

- Outer Process Sandbox
- Inner Static Verification Sandbox

# Process Sandbox

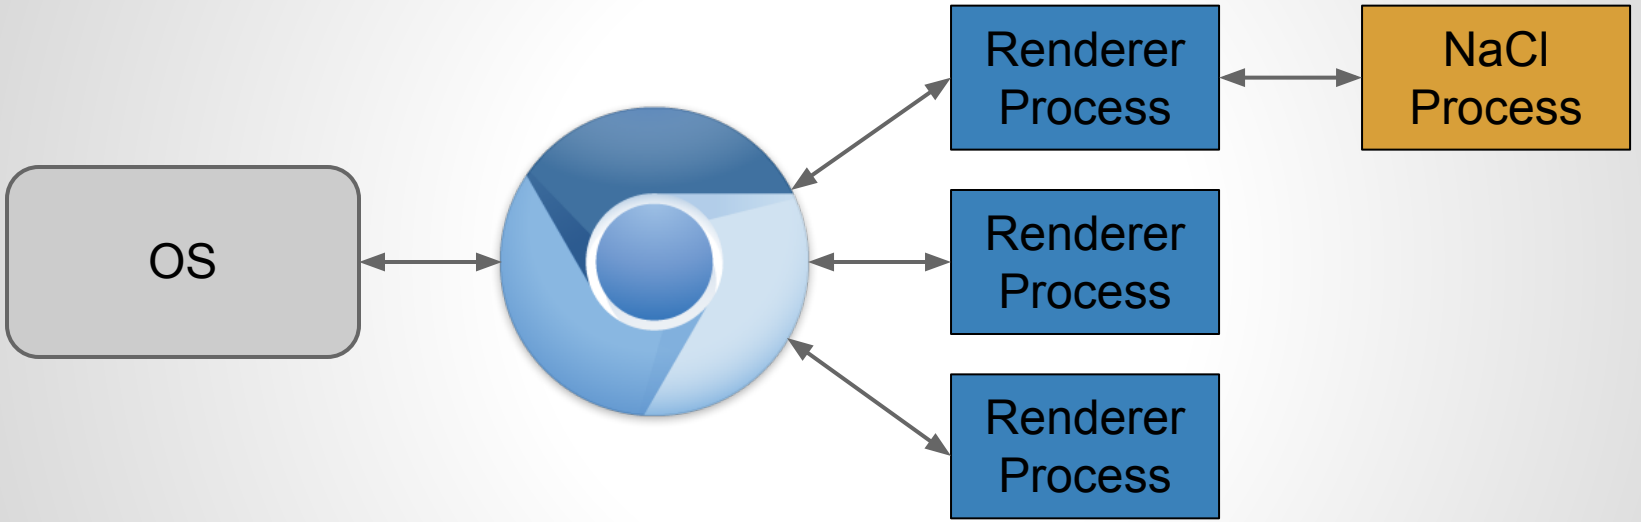

# Software Fault Isolation

- Restrict what code can execute (Control Flow Integrity → CFI)
- Restrict what data can be accessed (Data Integrity)
- Restricted I/O interface

# NaCl Address Space

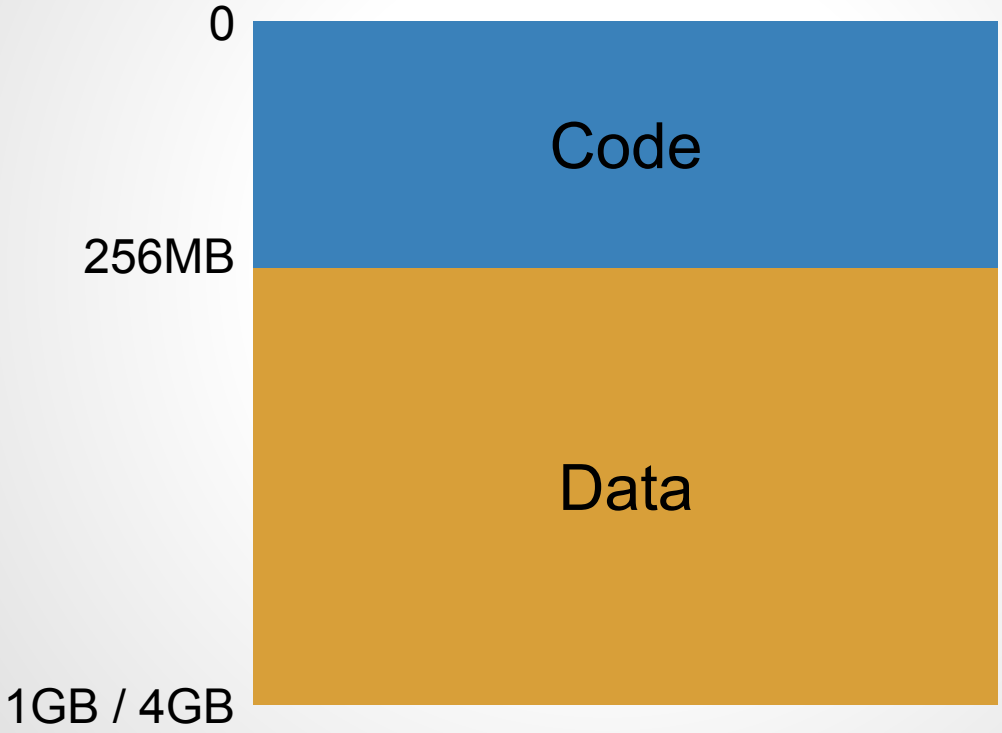

# Code Flow Integrity

- Be able to disassemble all code
- Disassembly must be unambiguous
- Know where all direct jumps can lead
- Know all indirect jumps (including returns) are safe

# Bundle Alignment

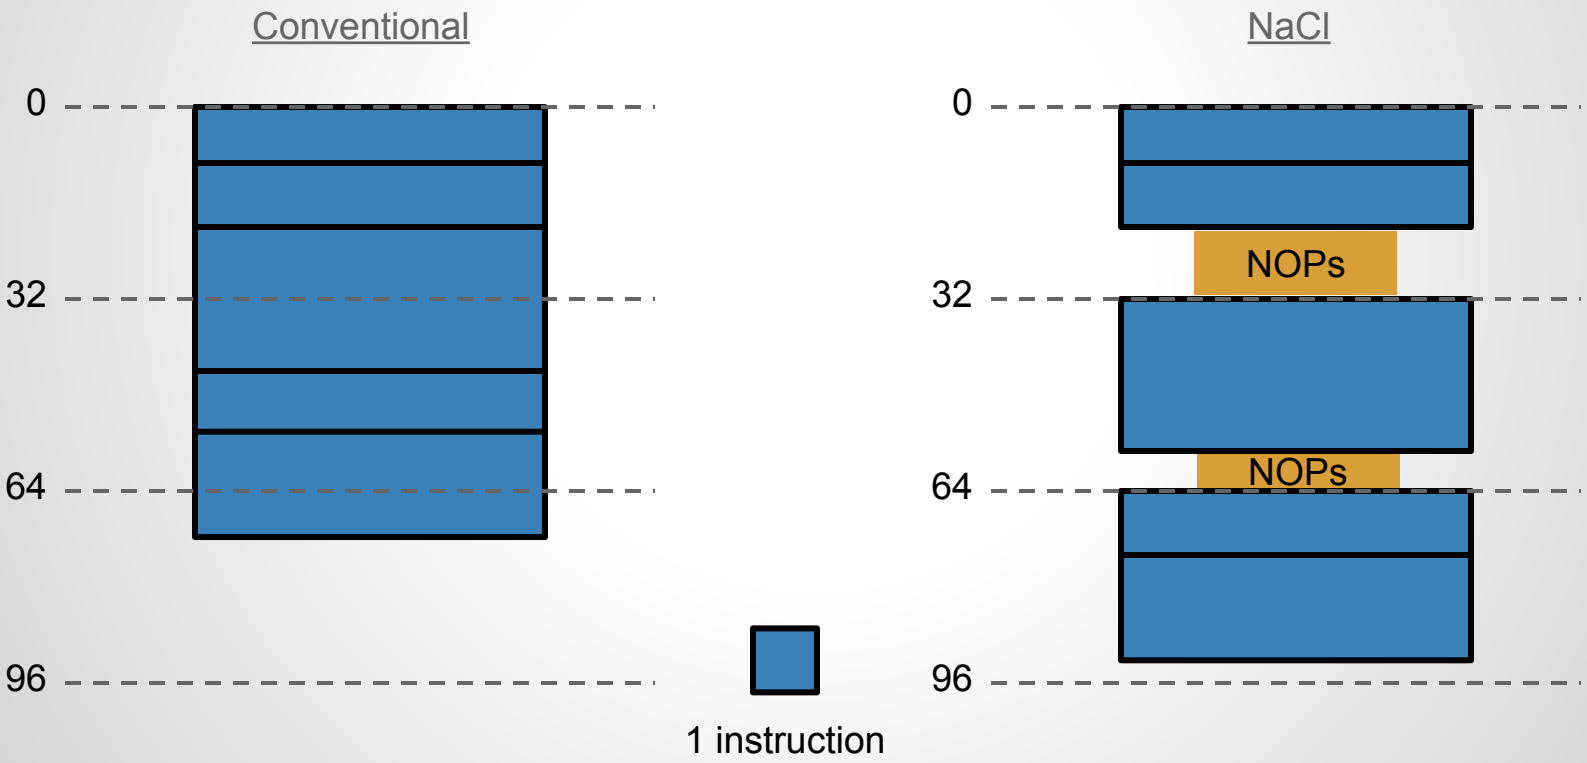

# Code Flow / Data Integrity

- Architecture specific restrictions on valid regions (segments)
- Mask jumps / memory access

# Syscalls

- Ban all the dangerous instructions
- Do syscalls through exit trampolines

# PPAPI

- Asynchronous API for:
  - Graphics
  - Sound
  - Network Access
  - Clipboard
  - Browser local files

# PPAPI


PPAPI

C++
instance_->PostMessage(
pp::Var(...));

air_mech_x86_64.nexe

NaCl runtime

Native Client Process

# Chrome Web Store

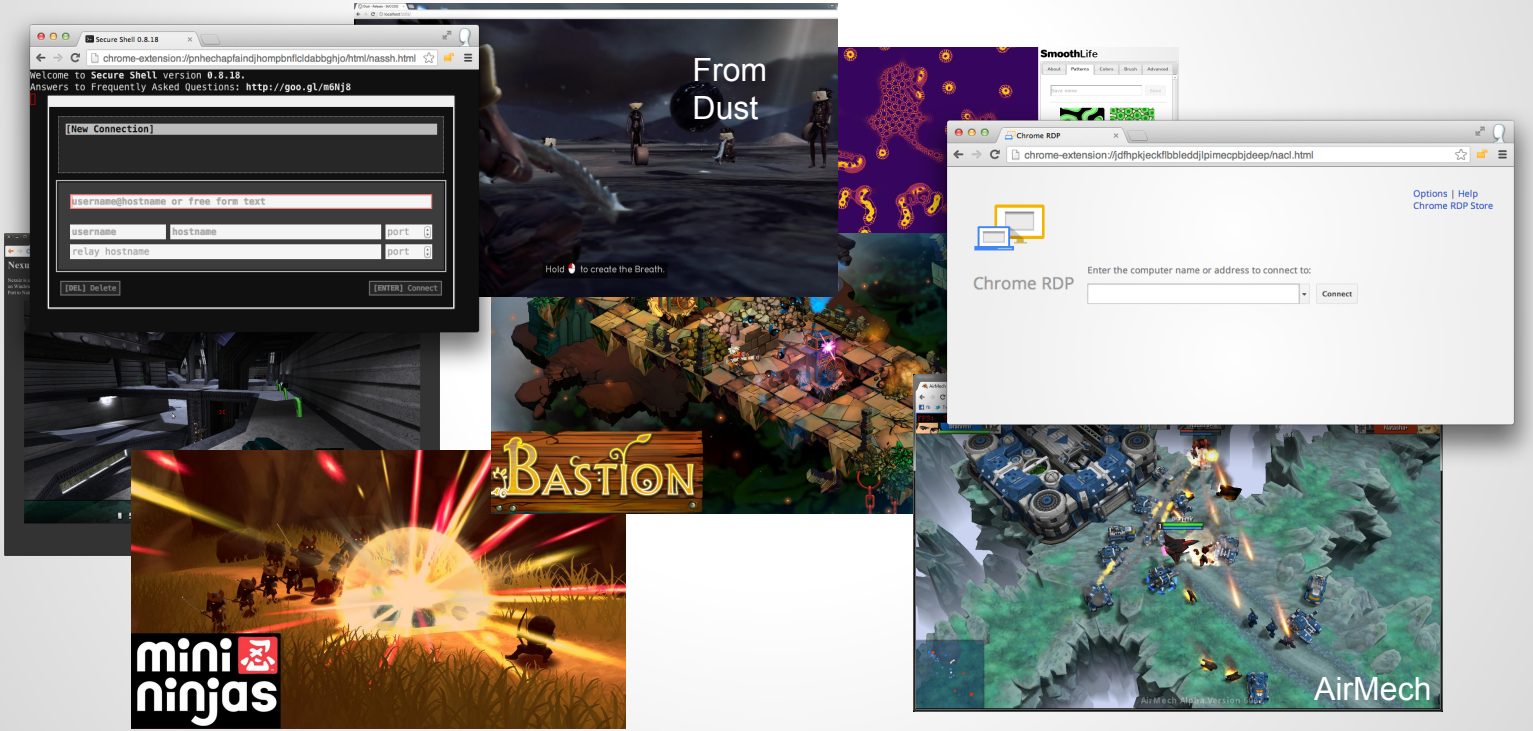

# Chrome Web Store

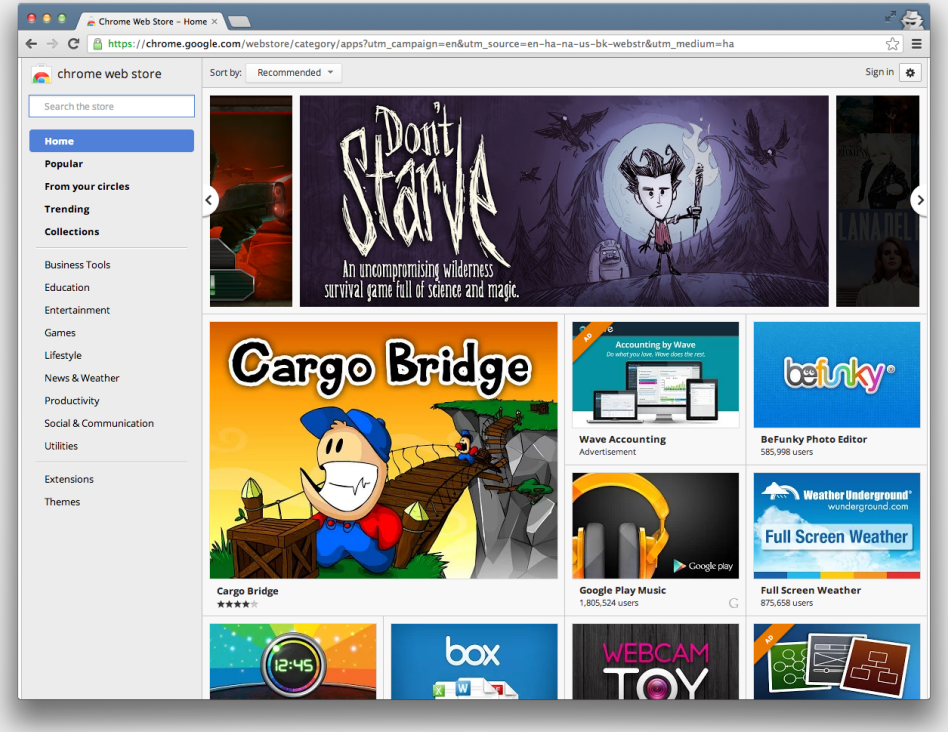

# nacl_io

- POSIX compatibility on top of PPAPI
- Virtual filesystem mounts:
  - Memory
  - DOM Storage
  - HTTP
  - FUSE
- POSIX like sockets
- Symbolic links (coming soon)

# nacl-spawn

- POSIX like processes
- PIDs
- spawnve, vfork, exec, waitpid
- Pipes (coming soon)

# nacl-spawn

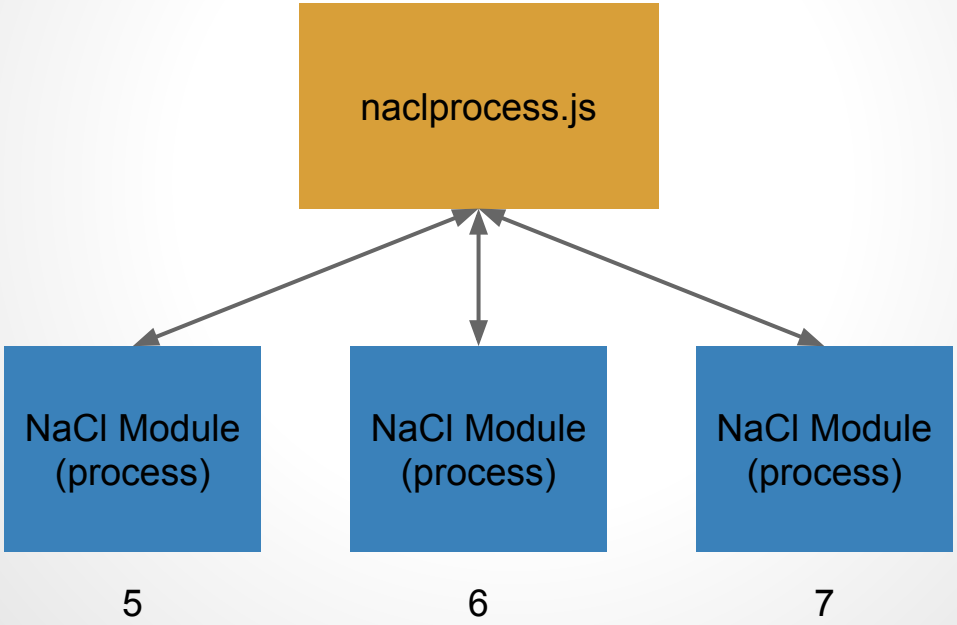
naclprocess.js
NaCl Module
(process)
NaCl Module
(process)
NaCl Module
(process)
5
6
7

# hterm

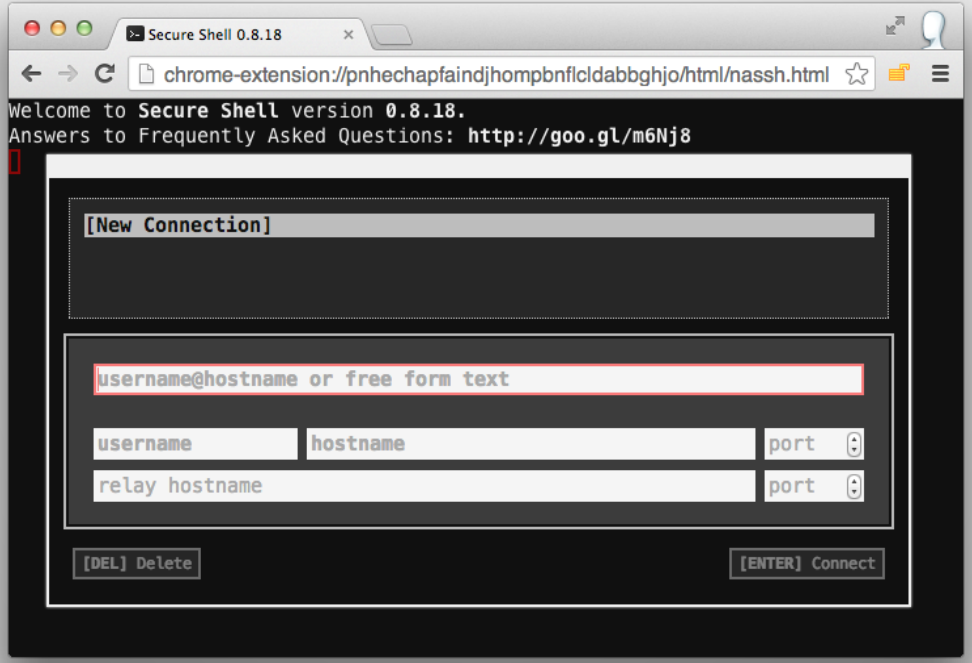
Secure Shell 0.8.18
chrome-extension://pnhechapfaindjhompbnflcldabbghjo/html/nassh.html
Welcome to Secure Shell version 0.8.18.
Answers to Frequently Asked Questions: http://goo.gl/m6Nj8
[New Connection]
username@hostname or free form text
username
hostname
port
relay hostname
port
[DEL] Delete
[ENTER] Connect

# Portability

- Most of what we want is C code
- But we want low level performance
- Move part of the compile to the client?

# Portable Native Client

- High level optimizations ahead of time (PNaCl Toolchain)
- Low level optimizations on the client (PNaCl Translator)
- Transmit portable type-reduced LLVM-IR (PEXE)

# The PNaCl Translator

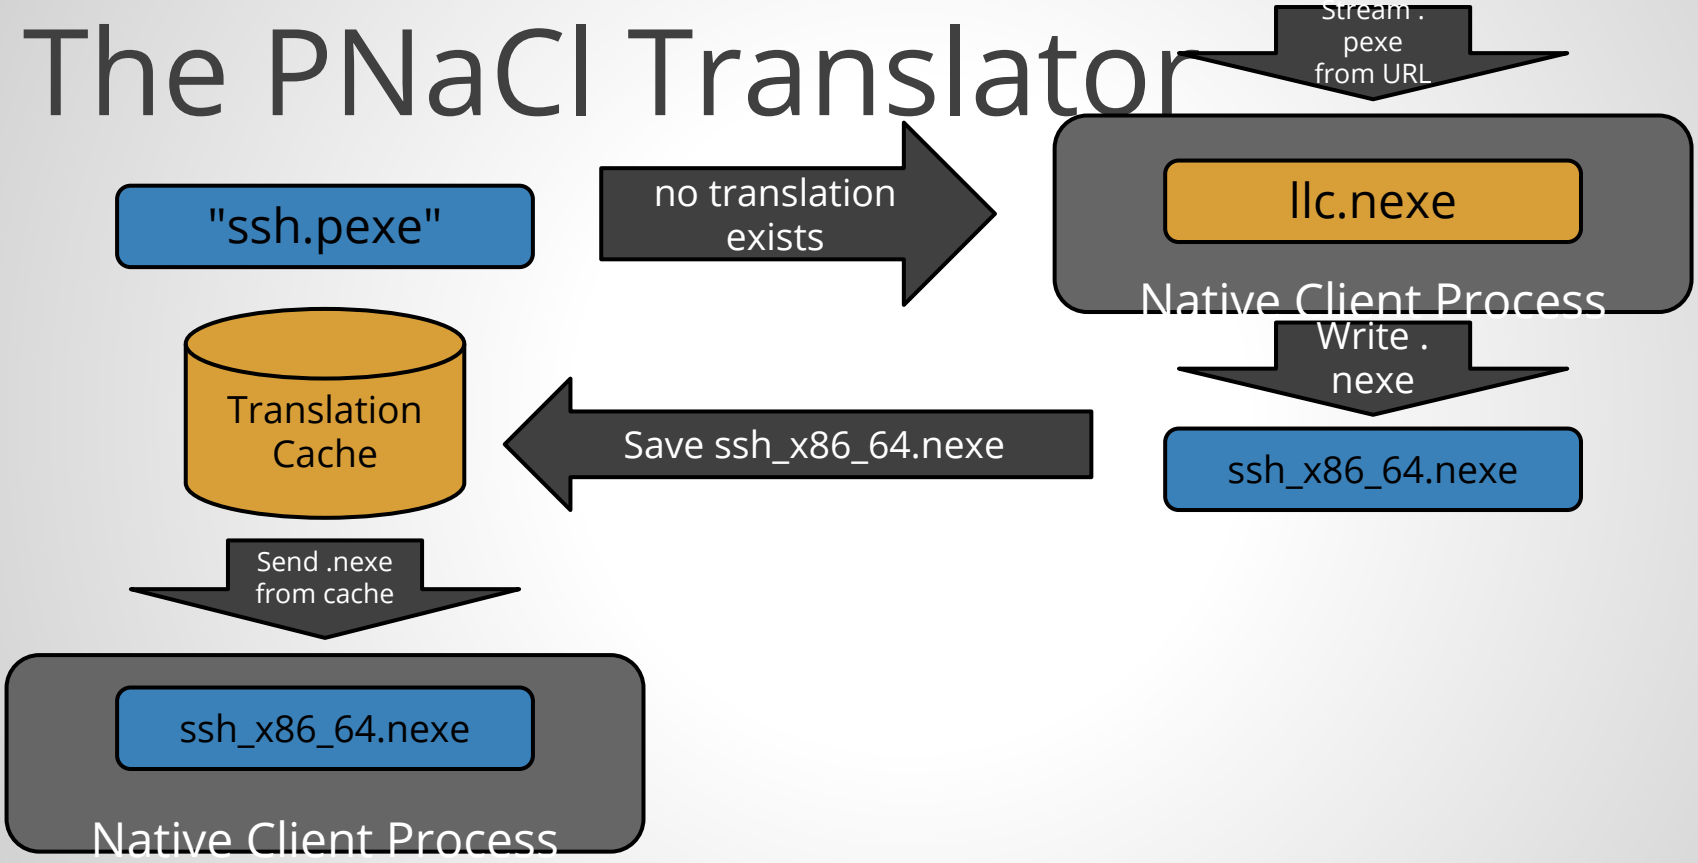

# PNaCl

- Custom compiler driver
- ABI simplification passes before and after optimization
- Translator running sandboxed in browser

# Naclports

- Free / Open Source ports to NaCl
- Stored as build scripts + patches
- Dependencies
- Upstream source mirroring
- Continuous build and test

# 200+ Ported Packages!

agg agg-demo alut-demo apr apr-util bash bdftopcf bigreqsproto binutils blackbox blis bochs boost boringssl box2d bullet busybox bzip2 cairo cfitsio civetweb clapack compositeproto coreutils curl damageproto devenv devenv-latest devil dosbox dreadthread drod emacs emacs-x expat faac faad2 ffmpeg fftw fftw-float fixesproto flac font-adobe-100dpi font-adobe-75dpi font-adobe-utopia-100dpi font-adobe-utopia-75dpi font-adobe-utopia-type1 font-alias font-bh-100dpi font-bh-75dpi font-bh-lucidatypewriter-100dpi font-bh-lucidatypewriter-75dpi font-bh-ttf font-bh-type1 font-bitstream-100dpi font-bitstream-75dpi font-bitstream-type1 font-cronyx-cyrillic font-cursor-misc font-daewoo-misc font-dec-misc font-ibm-type1 font-isas-misc font-jis-misc font-micro-misc font-misc-cyrillic font-misc-ethiopic font-misc-meltho font-misc-misc font-mutt-misc font-schumacher-misc font-screen-cyrillic font-sony-misc font-sun-misc font-util font-winitzki-cyrillic font-xfree86-type1 fontconfig fontsproto freealut freeimage freetype fvwm gc gcc gdb gettext geturl gforth giflib git glib glibc-compat glproto gmock gmp grep gsl gtest hdf5 icu imagemagick inputproto ipython-ppapi jpeg6b jpeg8d jsoncpp kbproto lame lcms leveldb libarchive libarchive-dev libatomic-ops libav libffi libfontenc libgit2 libgit2-demo libhangul libice libiconv libmikmod libmng libmodplug libogg liboggz libpciaccess libpng libpng12 libsm libsodium libssh libssh2 libtar libtheora libtomcrypt libtommath libtool libunistring libuuid libvorbis libwebp libx11 libxau libxaw libxcb libxext libxfont libxi libxinerama libxkbfile libxml2 libxmu libxpm libxrender libxt libxtst libyuv lua-ppapi lua5.1 lua5.2 m4 make mesa mesa-demo metakit mingn mongoose mp4v2 mpc mpfr mpg123 nacl-spawn nano ncurses netcat nethack ninja openal-ogg-demo openal-soft opencv openjpeg openscenegraph openssh openssl pango pcre physfs pixman protobuf protobuf-c protobuf25 python python-host python-ppapi python-static python3 python3-ppapi python_modules quakespasm randrproto readline recordproto regal renderproto resourceproto ruby ruby-ppapi scrnsaverproto scummvm sdl sdl-image sdl-mixer sdl-net sdl-tests sdl-ttf sdl2 sdl2-gfx sdl2-image sdl2-tests sdl2-ttf sdlquake snes9x speex sqlite subversion tar tcl texlive thttpd tiff tinyxml tk toybox twm unzip videoproto vim x264 xaos xbitmaps xcb-proto xcb-util xcmiscproto xextproto xeyes xfonts xineramaproto xkbcomp xkeyboard-config xorg-server xproto xtrans yajl zeromq zlib

# Continuous Build

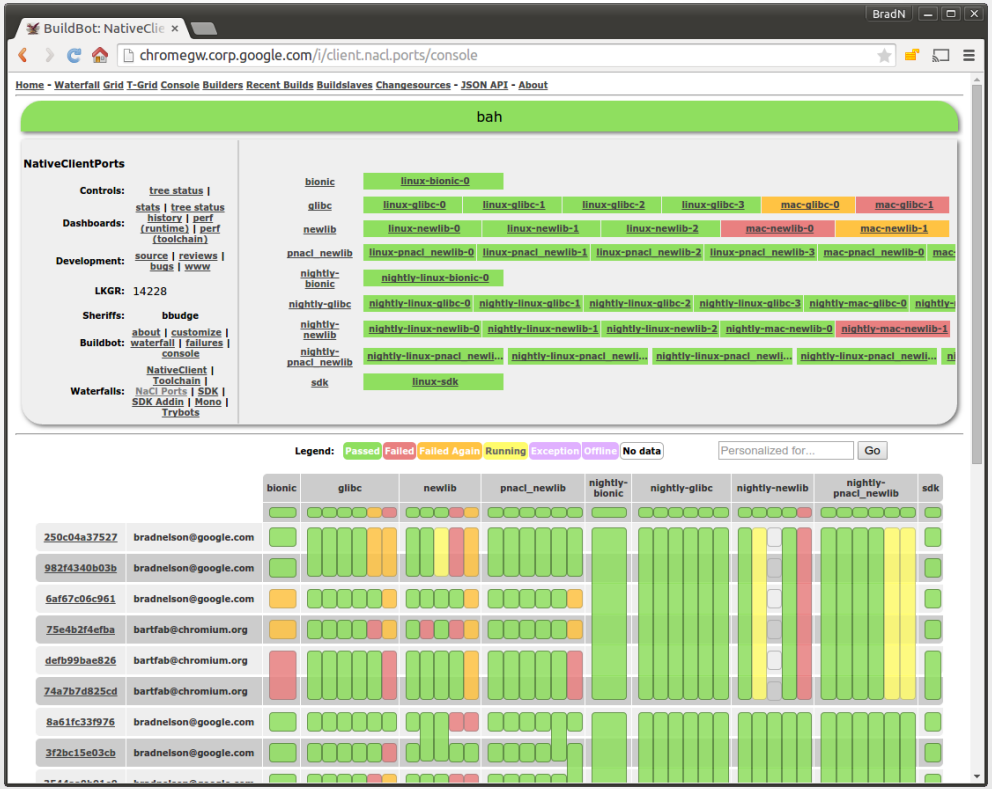

# NaCl Dev Environment

- Mini-Unix
- Bash
- Editors (Vim, Nano, Emacs)
- Interpreters (Python, Ruby, Lua, Tcl/Tk, GForth)
- Compilers (GCC, LLVM)

# NaCl Dev Environment

# PNaCl Toolchain

- LLVM
- Clang
- Binutils
- Libraries

# Porting Challenges

- Newlib LibC
  - glibc-compat
- Static Linking only
- Gold Linker plugin
- Compiler driver in Bash + Python
- Change main to nacl_main

# Static Linking

- Linker plugin (as used with LTO)
- We only need one plugin
- Modify gold to reference a single statically linked plugin
- Link LLVM + gold

# Compiler Drivers

- Bash + Python Compiler Driver
- Use nacl-spawn to create subprocess in:
  - Bash
  - Python
  - Clang

# Results

- 47 MB PEXE → 54MB NEXE (opt)
- 15 secs to compile a small sample
  - Why so slow?
  - DOM fs
- Can build fire example in NaCl Dev Env

# Possible Improvements

- Speed up DOM file system
- Replace DOM file system with ISO for headers
- Ditch our Python driver scripts

# Future Directions

- In browser development
- Cross-compile to native

# Thanks to awesome folks

- Sam Clegg (naclports + nacl_io)
- Ben Smith (naclports + nacl_io)
- Shinichiro Hamaji (gcc port + nacl-spawn)
- Pete Williamson (emacs port)
- Robert Ginda (hterm)
- Sergii Pylypenko (xserver on sdl)
- Austin Benson, Torin Rudeen, Channing Huang (awesome interns)
- Many more...

# Demo + Questions

http://gonacl.com/fire