

# Network stack virtualization

Marko Zec

[zec@fer.hr](mailto:zec@fer.hr)

# *Progress since Cambridge (August)*

---

- Diff against head reduced from ~1M to 280k
  - 2 big / intrusive commits to head, introducing:
    - V\_ macros, currently resolving to original globals
      - `s/ipforwarding/V_ipforwarding/`
      - `#define V_ipforwarding ipforwarding`
    - Context-setting macros, resolving to whitespace
      - `INIT_VNET_INET(curvnet)`
    - `SYSCTL_*(..)`  $\rightarrow$  `SYSCTL_V_*(V_NET, vnet_xxx, ...)`
    - bz@ proved all (sic) changes so far have no functional impact

# *Progress since Cambridge (August)*

---

- 1+ month behind schedule, step 1 and 1.5 done:
  - <http://wiki.freebsd.org/Image/Notes200808DevSummit>
- Rethinking step #2 – split in two parts
  - 2.1: implement constructor functions
    - `-int ipforwarding = 0;`
    - `+int ipforwarding;`
    - `@ip_init():`
    - `+ V_ipforwarding = 0;`
  - 2.2: only then (conditionally) move globals to structs
    - Globals: `#ifndef VIMAGE_USE_STRUCTS`

## *Original commit plan step #2*

---

- Add changes to convert all globals to members of per-module structures. Macros defined so that depending on compile options structures or globals are used (one global structure).
- Performance implications of using structures are evaluated. Structures possibly tuned.
- Initialisation routines added, checked and tuned.
  - `#if VIMAGE_USE_STRUCTS`
  - `#define V_hostname sys_globals.hostname`
  - `#else`
  - `#define V_hostname hostname`
  - `#endif`

## *Commit plan step #3*

---

- Add destructor functions.
- Add / activate iterator macros.
- Activate context pointers.
  - `INIT_VNET_INET()` would start doing something.
- Better documentation / manual pages
  - /me: Prepare in advance

## *Commit plan step #4*

---

- JAIL+Vimage framework committed. e.g. add new syscall, program, etc. (part one of meat diffs) structures still only global instances.
- vimage enhanced jails can be created but act just like normal jails?
- Ability to create > 1 vimage enabled. Vimage enhanced jails now have private network stacks etc.

# Perforce

---

- //depot/projects/vimage/...
  - Tracking HEAD
  - Ad-hoc design / porting document draft there:  
**porting\_to\_vimage.txt**
- //depot/projects/vimage-commit2/...
  - THE branch to eventually get committed to svn
  - Initially machine generated changes to HEAD (var\_rename.tcl), then manually adjusted
  - Currently identical to HEAD

# *jail\_set\_vimage*

---

- Perforce:
  - //depot/user/jamie/jail\_set\_vimage
    - Jamie Gritton's merge of vimage and jail\_set hacks
    - Basically, vimage framework reduced to network stack virtualization, hanging off “new jails”
    - [http://gritton.org/jail\\_set\\_vimage.diff](http://gritton.org/jail_set_vimage.diff) - diff against //depot/projects/vimage/