

Zero to Production

Erlang Factory

San Francisco

March 11, 2016

Susan Potter @ Lookout

twitter: @SusanPotter

github: mbbx6spp

InfraEng @ Lookout

```
1 # finger infraeng
2 Login: infraeng
3 Name: Infra Eng @ Lookout
4 Shell: /run/current-system/sw/bin/bash
5 Last login Mon Mar 11 14:10 (PST) on pts/10
6
7 * Multiple services in prod
8 * 200-300 hosts monitored already
9 * Internal Nix channel
10 * Internal binary cache
11 * One repository per service
12 * Repository is source of truth
13 * We are hiring! Come talk to me. :)
```

% whoami

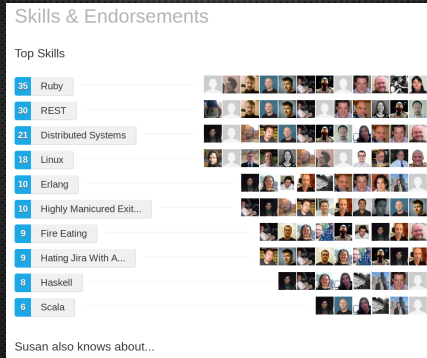


Figure: From backend dev to infrastructure engineering

Reliability

“Those who want really reliable software will discover that they must find means of avoiding the majority of bugs to start with, and as a result the programming process will become cheaper.” – EWD340



Reduce Costs & Frustration

*“If you want more effective programmers, you will discover that they should not waste their time debugging, they should not introduce the bugs to start with.”–
EWD340*

Why care now?

- 1 Economic factors
large distributed deployments

Why care now?

- ① **Economic factors**
large distributed deployments
- ② **Human factors**
high churn/turnover, low quality of ops life

Why care now?

- ① **Economic factors**
large distributed deployments
- ② **Human factors**
high churn/turnover, low quality of ops life
- ③ **Technological factors**
programmable infrastructure & FP no longer just for academics

More Services

- 1 Currently 20-30 services
- 2 More services ready each month
- 3 Expect 50+ by end of year
- 4 Various stacks/runtimes

More Environments

- ① Ephemeral (integration testing)
- ② Product lines (consumer vs enterprise)
- ③ Performance
- ④ Partners

More Persistence



redis



MySQL®

elastic



hadoop

STORM



cassandra



Agenda

- 1 The Problem
- 2 The Principle
- 3 Introduce Nix* Ecosystem
- 4 How Nix Solves Our Problems
- 5 Lessons Learned

Problem: Software Delivery

Environment provisioning not repeatable in practice

Problem: Software Delivery

Continuous integration builds
break with app dependency
changes

Problem: Software Delivery

Deploys have unexpected
consequences that
'-dry-run/-why-run' cannot
catch

Requirements: Optimize for ...

- 1 Scalability solved by on-demand "cloud"

Requirements: Optimize for ...

- 1 Scalability solved by on-demand "cloud"
- 2 Reliability solved by ...

So what yields reliability?

So what yields reliability?

Ability to reason about code.

What allows you to reason about code?

What allows you to reason about code?

Referential transparency (RT)!

RT Refresher

Functions have inputs (Erlang)

```
1 -module(myfuncs).  
2  
3 % Two input arguments here  
4 myadd(X, Y) -> X + Y.  
5  
6 % One input argument here  
7 mylen(S) -> len(S).
```


Functions have inputs (Nix)

```
1 # stdenv, fetchurl, gcc, help2man are
2 # (package) inputs to this package
3 { stdenv, fetchurl, gcc, help2man }:
4 let
5     version = "2.1.1";
6 in stdenv.mkDerivation {
7     inherit version;
8     name = "hello-${version}";
9     src = fetchurl { ... };
10    # gcc and help2man are build deps
11    buildInputs = [ gcc help2man ];
12 }
```

Functions return a result (Erlang)

```
1 Eshell V7.0 (abort with ^G)
2 1> c(myfuns).
3 {ok,myfuns}
4 2> myfuns:myadd(1,4).
5 5
6 3> myfuns:mylen("Hello, Erlang Factory.").
7 22
8 4> q().
9 ok
10 5>
```

Functions return a result (Nix)

```
1 $ nix-repl '<nixpkgs>'
2
3 Loading <nixpkgs>\ldots
4 Added 5876 variables.\
5
6 nix-repl> hello = import ./hello.nix { \
7   inherit stdenv fetchurl gcc; \
8 }
9
10 nix-repl> hello
11 derivation /nix/store/...0am-hello-2.1.1.drv\
```

Functions return a result (Nix)

```
1 nix-repl> "${hello}"
2 "/nix/store/jg1l1...lsj-hello-2.1.1"
3
4 nix-repl> :q
5
6 $ nix-build hello.nix \
7   --arg stdenv "(import <nixpkgs> {}).stdenv" \
8   --arg fetchurl "(import <nixpkgs> {}).fetchurl" \
9   --arg gcc "(import <nixpkgs> {}).gcc" \
10  --arg help2man "(import <nixpkgs> {}).help2man"
11 /nix/store/jg1l1...lsj-hello-2.1.1
```

Only depend on inputs (Erlang)

```
1 \ $ cat inputs.erl
2 -module(inputs).
3
4 mylol(X, Y) -> Z.
5 \ $ erl
6 Eshell V7.0 (abort with ^G)
7 1> c(inputs).
8 inputs.erl:3: variable 'Z' is unbound
9 inputs.erl:3: Warning: function mylol/2 is unused
10 inputs.erl:3: Warning: variable 'X' is unused
11 inputs.erl:3: Warning: variable 'Y' is unused
12 error
13 2> q().
14 ok
```

Only depend on inputs

```
1 # Remove help2man from package input arguments
2 \$$ cat hello.nix hello.nix.1
3 1c1
4 < { stdenv, fetchurl, gcc, help2man }:
5 ---
6 > { stdenv, fetchurl, gcc }:
7
8 \$$ nix-build hello.nix.1 \
9   --arg stdenv "(import <nixpkgs> {}).stdenv" \
10  --arg fetchurl "(import <nixpkgs> {}).fetchurl" \
11  --arg gcc "(import <nixpkgs> {}).gcc"
12 error: undefined variable help2man at hello.nix:11:23
```

Only depend on inputs

```
1 # Remove help2man from buildInputs
2 \$$ cat hello.nix hello.nix.2
3 11c11
4 < buildInputs = [ gcc help2man ];
5 ---
6 > buildInputs = [ gcc ];
7
8 \$$ nix-build hello.nix.2 \
9 --arg stdenv "(import <nixpkgs> {}).stdenv" \
10 --arg fetchurl "(import <nixpkgs> {}).fetchurl" \
11 --arg gcc "(import <nixpkgs> {}).gcc" \
12 --arg help2man "(import <nixpkgs> {}).help2man"
13 ...
```

Only depend on inputs

```
1 these derivations will be built:
2   /nix/store/19x32rhqx...mn80-hello-2.1.1.drv
3 building path(s) /nix/store/v38...2m58h-hello-2.1.1
4 unpacking sources
5 unpacking source archive /nix/store/...-hello-2.1.1.tar.gz
6 source root is hello-2.1.1
7 ...
8 /nix/...-bash-.../bash: help2man: command not found
9 Makefile:282: recipe for target 'hello.1' failed
10 make[2]: *** [hello.1] Error 127
11 ...
12 error: build of /nix/...n80-hello-2.1.1.drv failed
```


Return same result given same inputs

```
1 prop_ref_trans() ->
2   ?FORALL({X, Y}, {integer(), integer()}),
3     begin
4       Z0 = myadd(X, Y),
5       Z1 = myadd(X, Y),
6       Z0 ::= Z1
7     end).
```

Return same result given same inputs

```
1 $ while true; do
2   nix-build \
3     --arg stdenv "(import <nixpkgs> {}).stdenv" \
4     --arg fetchurl "(import <nixpkgs> {}).fetchurl" \
5     --arg gcc "(import <nixpkgs> {}).gcc" \
6     --arg help2man "(import <nixpkgs> {}).help2man" \
7     hello.nix
8 done
9 /nix/store/jg1l1kw...sj-hello-2.1.1
10 /nix/store/jg1l1kw...sj-hello-2.1.1
11 ...
12 /nix/store/jg1l1kw...sj-hello-2.1.1
13 ^Cerror: interrupted by the user
```

The Big idea

Referential Transparency

Given same inputs, return same result. Always.

Questions so far?



Figure: Awake?

Mainstream Package Management

Based on shared + mutable state (filesystem)

Violates RT



Alternative Approaches

- shared + immutable
- private + mutable
- expensive coarse grained locks
- hybrid without the expense

Define all inputs

- Force clean build env (chroot)
- Requires explicit inputs
- Full dependency definition

Ensure RT

- Use private mutable space
- Different inputs, different result
- Symlink unique results (atomic op)

Nix Ecosystem

- Expression language: Nix
- Package management: Nix
- Channel: 'nixpkgs'
- Operating System: NixOS
- Configuration "modules": NixOS modules
- Provisioning: NixOps
- Orchestration: Disnix
- CI: Hydra

Repeatable Dev Envs

```
1 $ nix-shell -p erlangR17_odbc
2 these paths will be fetched (37.65 MiB download, 112.65 MiB unpacking)
3   /nix/store/0jvs...3vd-unixODBC-2.3.2
4   /nix/store/wf7w...6fp-erlang-17.5-odbc
5 fetching path /nix/store/0jvs...-unixODBC-2.3.2...
6 ...
7 [nix-shell:~]$ erl
8 Erlang/OTP 17 [erts-6.4] [source] [64-bit] ...
9 Eshell V6.4 (abort with ^G)
10 1>
```

Repeatable Dev Envs

```
1 $ nix-shell -p erlangR18_javac
2 these paths will be fetched (38.04 MiB download, 113.9 MiB unpacked)
3   /nix/store/94a...b3xn-erlang-18.2
4 fetching path /nix/store/94a...b3xn-erlang-18.2...
5 ...
6 [nix-shell:~]$ erl
7 Erlang/OTP 18 [erts-7.2] [source] [64-bit] ...
8
9 Eshell V7.2  (abort with ^G)
10
11 1>
```

Repeatable Dev Envs

```
1 $ declare pkgghost="releases.nixos.org"
2 $ declare release_url="https://${pkgghost}/nixos"
3 $ nix-channel --add \
4   "${release_url}/16.03-beta/nixos-16.03.30.2068621" nixpkgs
5 $ nix-shell
6 these derivations will be built:
7   /nix/store/267y...-elm-0.16.0.drv
8 these paths will be fetched (31.49 MiB download, 379.9 MiB unpack)
9   /nix/store/0bkd...-scientific-0.3.4.4
10  /nix/store/0d3y...-nodejs-4.3.1
11  ...
12 building path(s) /nix/store/jjzr...-elm-0.16.0
13 created 6 symlinks in user environment
```

Repeatable Dev Envs

```
1 $ cat shell.nix
2 { pkgs ? import <nixpkgs> {}, ... }:
3 let
4   inherit (pkgs) stdenv;
5 in stdenv.mkDerivation {
6   name = "myerlprj-devenv";
7   buildInputs = with pkgs; [
8     gitFull           # Developer dependency
9     erlangR18         # Erlang version to use
10    hex2nix rebar3    # Erlang dev cycle tools
11    postgresql        # RDBMS
12    elmPackages.elm   # for front-end compiler
13  ];
14  ...
```

Repeatable Dev Envs

```
1  ...
2  shellHook = ''
3      export SERVICE_PORT=4444
4      export DATABASE_PORT=5432
5      export DATABASE_PATH=$PWD/data
6      export LOG_PATH=$PWD/log
7      if [ ! -d "${DATABASE_PATH}" ]; then
8          initdb "${DATABASE_PATH}"
9      fi
10     pg_ctl -D "${DATABASE_PATH}" \
11         -l "${LOG_PATH}" \
12         -o --path="${DATABASE_PATH}" start
13 '';
14 }
```

Consistent CI Deps

```
1 $ head -3 z/ci/verify
2 #!/usr/bin/env nix-shell
3 #!nix-shell -I nixpkgs=URL
4 #!nix-shell -p erlangR18 postgresql -i bash
```


Consistent CI Deps

```
1 ...
2 set -eu
3
4 ! test -d "${DATABASE_PATH}" && \
5     initdb "${DATABASE_PATH}"
6 elm-make priv/elm/*
7 rebar3 clean compile dialyzer
8 pg_ctl -D "${DATABASE_PATH}" \
9     -l "${LOG_PATH}" -o \
10     --port="${DATABASE_PORT}" start
11 rebar3 ct
12 pg_ctl -D "${DATABASE_PATH}" stop
```

Consistent CI Deps

- Pin channel versions = source + CI consistency
- Update CI build deps with app code
- No OOB 'converge'-ing CI build hosts!

Predictable Deploys

- Diff dependency path tree
- Test node configuration in VM
- Test NixOS module logic
- Security auditing

Diff Dependencies

```
1 $ nix-store -qR /nix/store/*-myerlprj-*
2 /nix/store/8jhy2j7v0mpwybw13nd4fjlsfqc9xnlh-write-mirror-11
3 /nix/store/17h0mw5sipbvg70hdsn8i5mai4619l8c-move-docs.1h
4 ...
5 /nix/store/p6gn7inwvm61phqw3whhlbl20n8c5dgb-git-2.7.1.driv
6 /nix/store/z2jvckzhy5322d9ir0xv2hbqp6yakayj-myerlprj-denv
```

Machine Config

```
1 { config, pkgs, ... }:
2 let
3   inherit (pkgs) lib;
4   ntpF = (idx: "${idx}.amazon.pool.ntp.org")
5   domain = "example.com";
6 in {
7   boot.cleanTmpDir = true;
8   boot.kernel.sysctl = {
9     "net.ipv4.tcp_keepalive_time" = 1500;
10    # other sysctl key-values here...
11  };
12  networking.hostName = "nixallthethings.${domain}";
13  networking.firewall.enable = true;
14  services.ntp.servers = map ntpF (lib.range 0 3);
15  services.zookeeper.enable = true;
16  security.pki.certificateFiles = [./internal_ca.crt];
17  time.timeZone = "UTC";
18 }
```

Test Machine Config (VM)

```
1 $ env NIXOS_CONFIG=$PRJROOT/priv/nix/config.nix \  
2   nixos-rebuild build-vm  
3 $ ./result/bin/run-hostname-vm  
4 ...  
5  
6 $ env NIXOS_CONFIG=$PRJROOT/priv/nix/config.nix \  
7   --target-host myerlprj-test-1.integ.bla \  
8   nixos-rebuild build-vm
```

Module Integration Testing

```
1 $ grep -A8 elasticsearch.enable $PWD/priv/nix/config.nix
2   elasticsearch.enable = true;
3   elasticsearch.jre =
4     mychannel.elasticsearch_2_2_0;
5   elasticsearch.jre =
6     mychannel.oraclejre8u74;
7   elasticsearch.node.name =
8     "elasticsearch-0.${domain}";
9   elasticsearch.dataDir =
10    [ "/data0" "/data1" "/data3" ];
```

Module Integration Testing

```
1 $ grep -A8 "node_health" $PWD/priv/nix/modules/elasticsearch
2 subtest "elasticsearch_node_health", sub {
3     $es0->waitForUnit("elasticsearch.service");
4     $es1->waitForUnit("elasticsearch.service");
5     $es0->succeed("${waitForTcpPort es0 9300 60}");
6     $es1->succeed("${waitForTcpPort es1 9300 60}");
7     $es0->succeed("${curl es0 9200 /}");
8     $es1->succeed("${curl es1 9200 /}");
9 }
```


Security Auditing

```
1 $ nix-store -qR /path/to/app/pkg | sort | uniq
2 /nix/store/002v...-libdc1394-2.2.3
3 /nix/store/04bw...-expat-2.1.0
4 /nix/store/04df...-haskell-x509-validation-ghc7.8.4-1.5.1-s
5 /nix/store/06p6...-packer-e3c2f01cb8d8f759c02bd3cfc9d27ccia
6 ...
7 /nix/store/zv9r...-perl-libwww-perl-6.05
8 /nix/store/zvgj...-pypy2.5-stevedore-0.15
9 /nix/store/zw00...-libpciaccess-0.13.3
10 /nix/store/zz78...-libdvbpsi-0.2.2
```

Security Auditing

```
1 $ nix-store -qR /run/current-system | grep openssl  
2 /nix/store/x1zwzk4hrvj5fz...9hyn-openssl-1.0.1p  
3 /nix/store/m4kzbwji9jkw71...lx92-openssl-1.0.1p
```

Tradeoffs

- **Provisioning not solved**
Nixops expressiveness vs Terraform 'coverage'
- **Steep learning curve**
Docs great reference, but bad for n00bs!
- **Lots of upfront setup**
Internal Nix channels vs nixpkgs fork curation

Benefits

- Repeatable dev envs
- Consistent CI
- Predictable deploys
- Real rollback

The Win!

Simplicity at its core

What Next?

- Nix AWS Provisioning
- Idris to Nix Backend
- NixBSD Anyone????
- Nix on BEAM???

Where to Next?

- **Nix Manual:**
<http://nixos.org/nix/manual>
- **NixOS Manual:**
<http://nixos.org/nixos/manual>
- **Nix Cookbook:**
<http://funops.co/nix-cookbook>
- **Nix Pills (by Lethalman)**

Questions




Figure: Heckle me @SusanPotter later too.