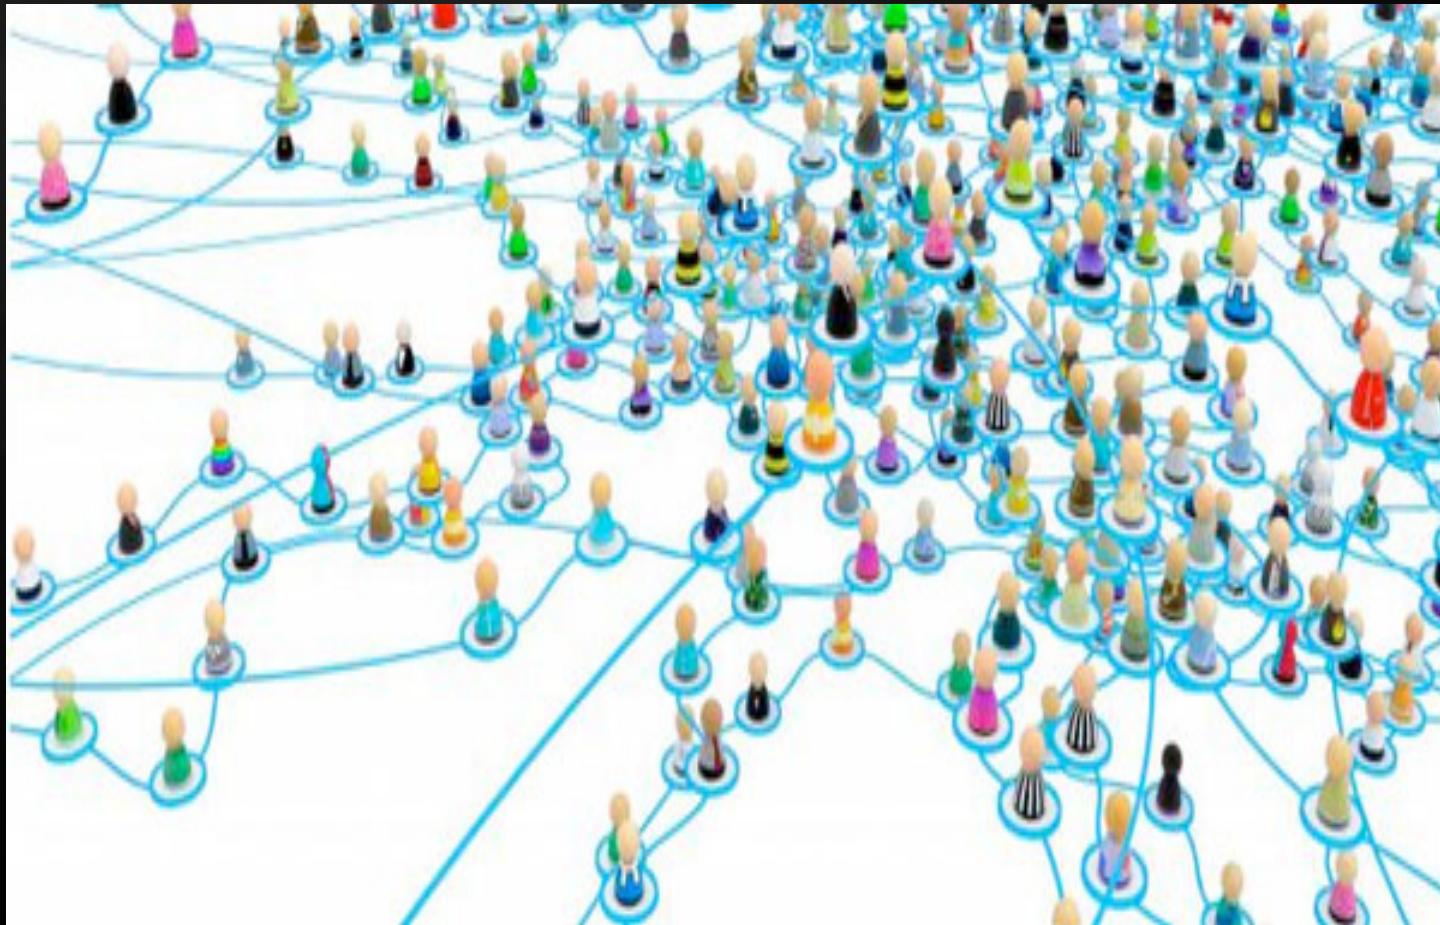


• Software Defined Networking, Erlang, and the Future of Computing



Stu Bailey, Founder/CTO

How Can We Program an Internet of Things?



How Can We Program a Million Cores?



And We Must Assume Hardware is Failing
(or changing) ALL the Time: “write once
run forever”

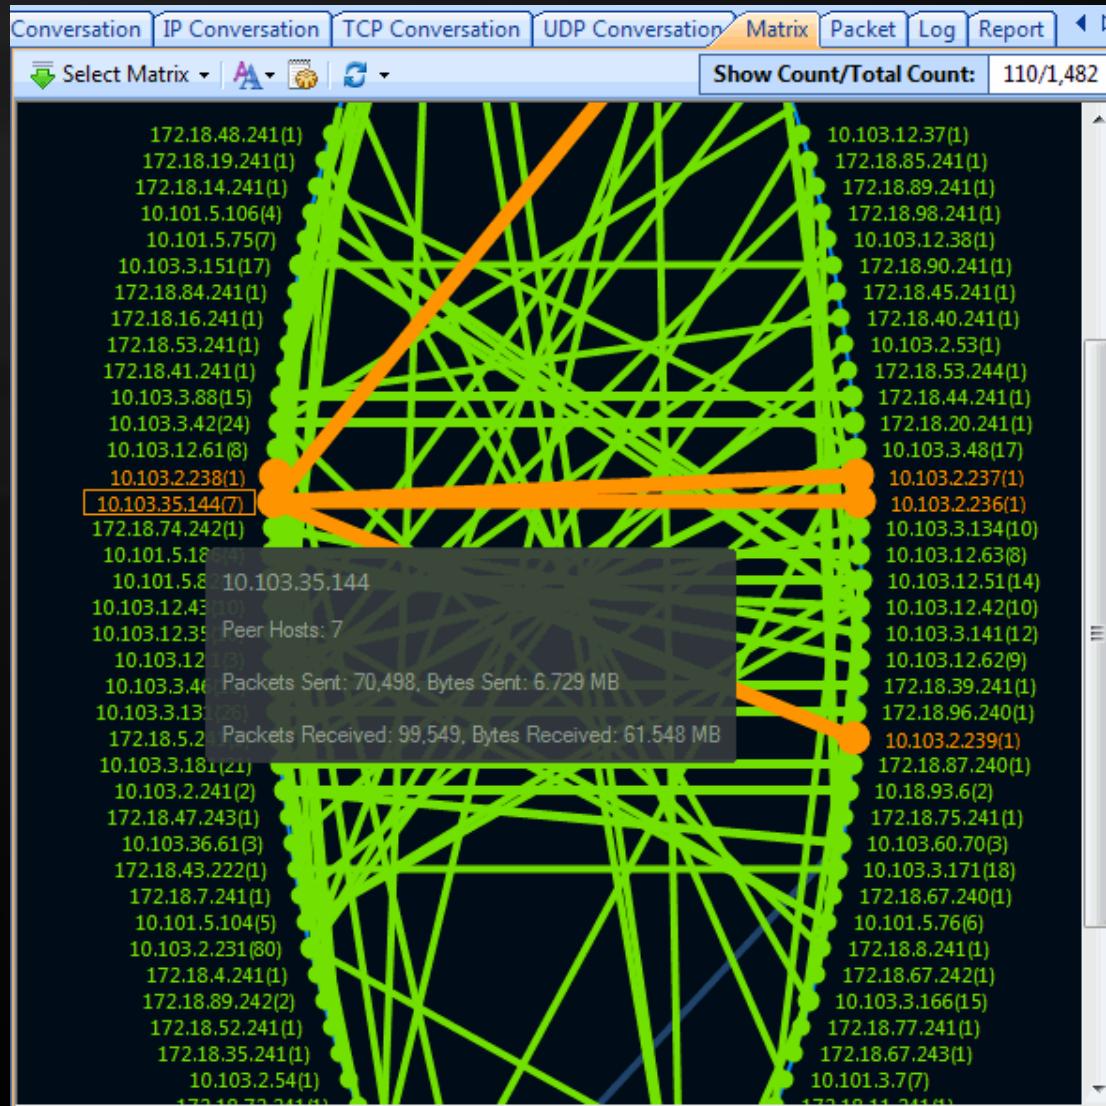


Of Course Erlang!!

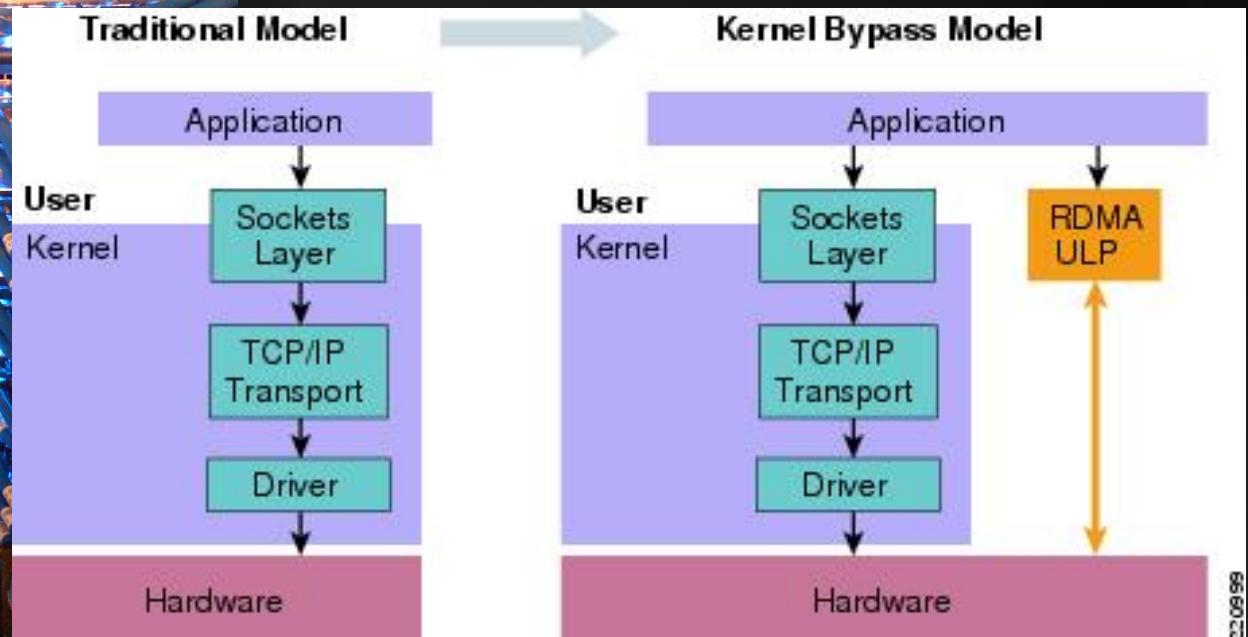
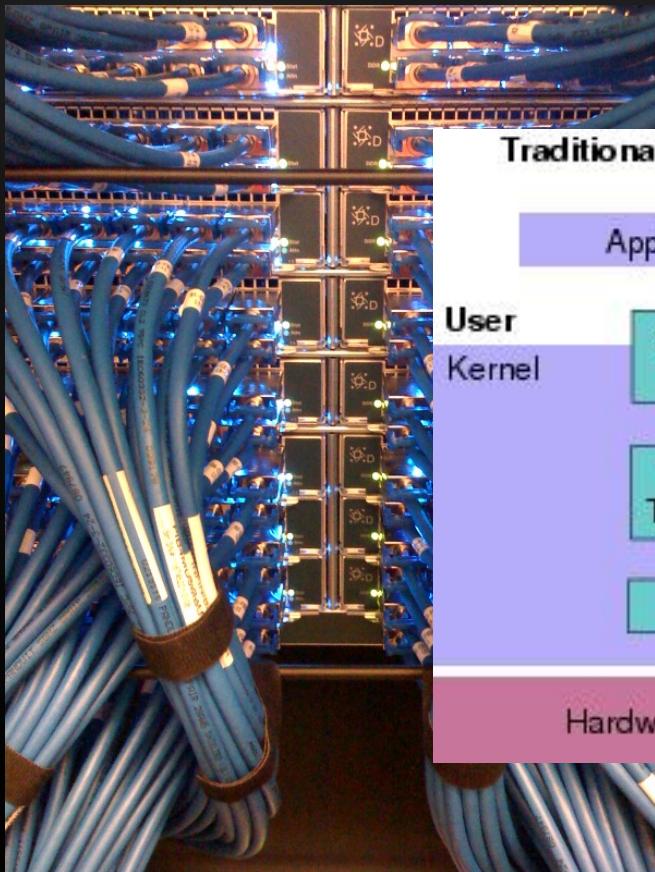


(or Scala/AKKA or node.js)

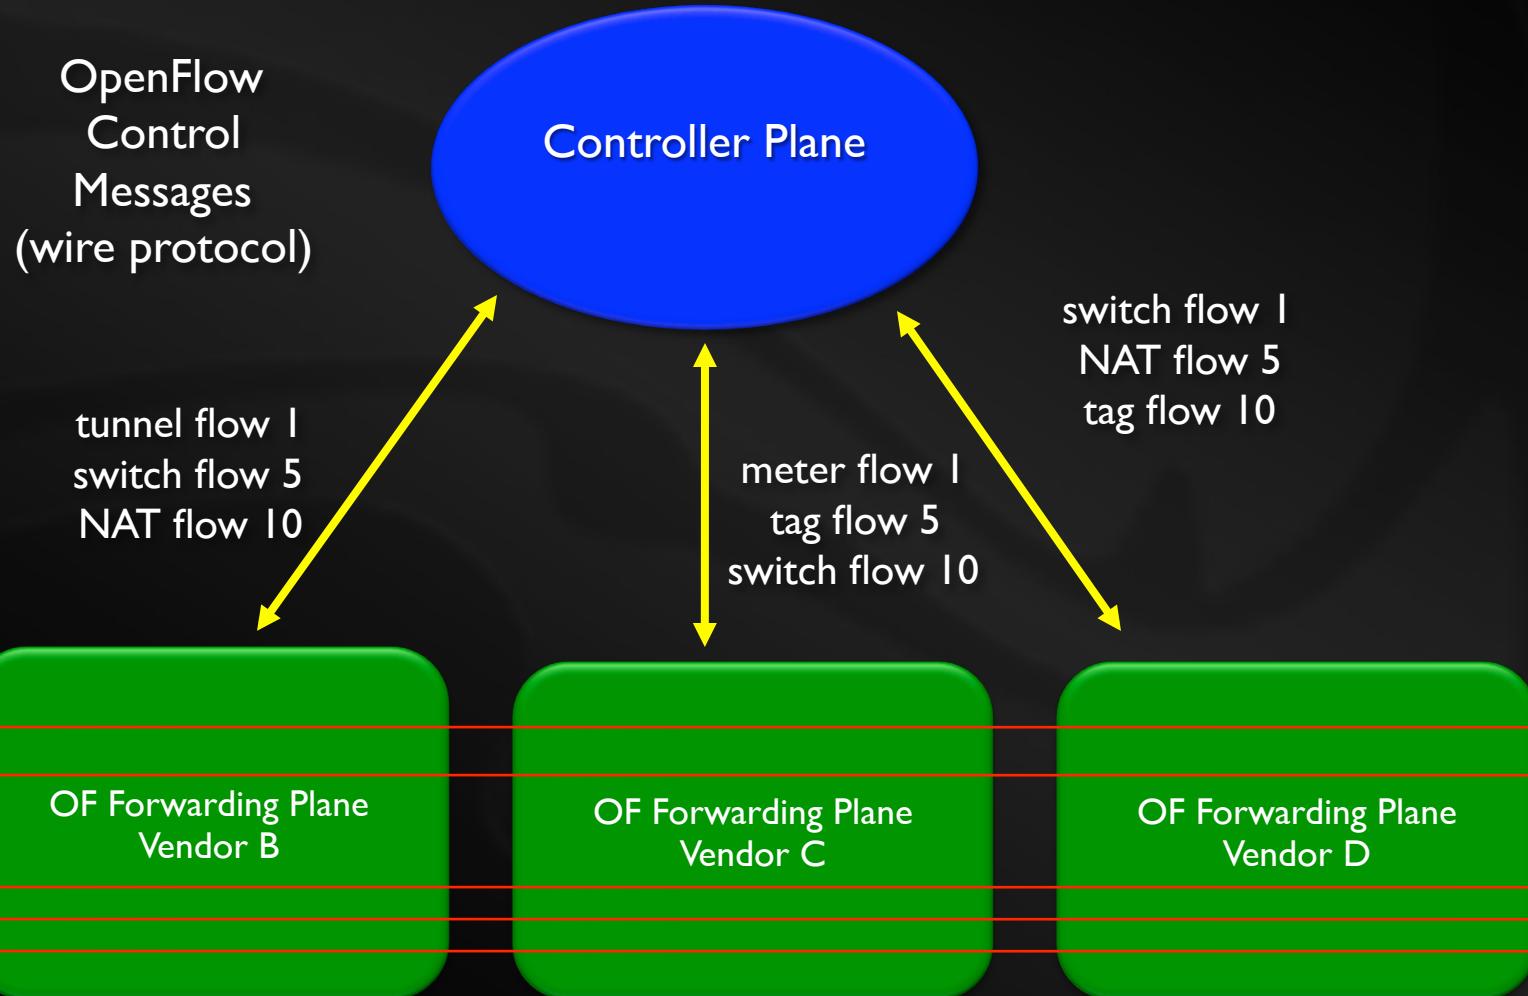
But wait! What about the network?!



Of course Infiniband! Except...



A STANDARD Forwarding Plane and Logically Centralized Controller Plane



Capabilities across forwarding plane vendors are fairly uniform
Performance and capacity are primary differentiators

FlowForwarding.org



OPEN SOURCE SDN STACK
ENTERPRISE FOCUSED

and maybe more?

Full SDN Stack Timeline

FlowForwarding
Community
launched
(June 2012)
+ alpha LINC Switch code availability

LINC Switch beta
v1.0
(Fall 2012)

LINC Switch beta
v2.X
(2013)

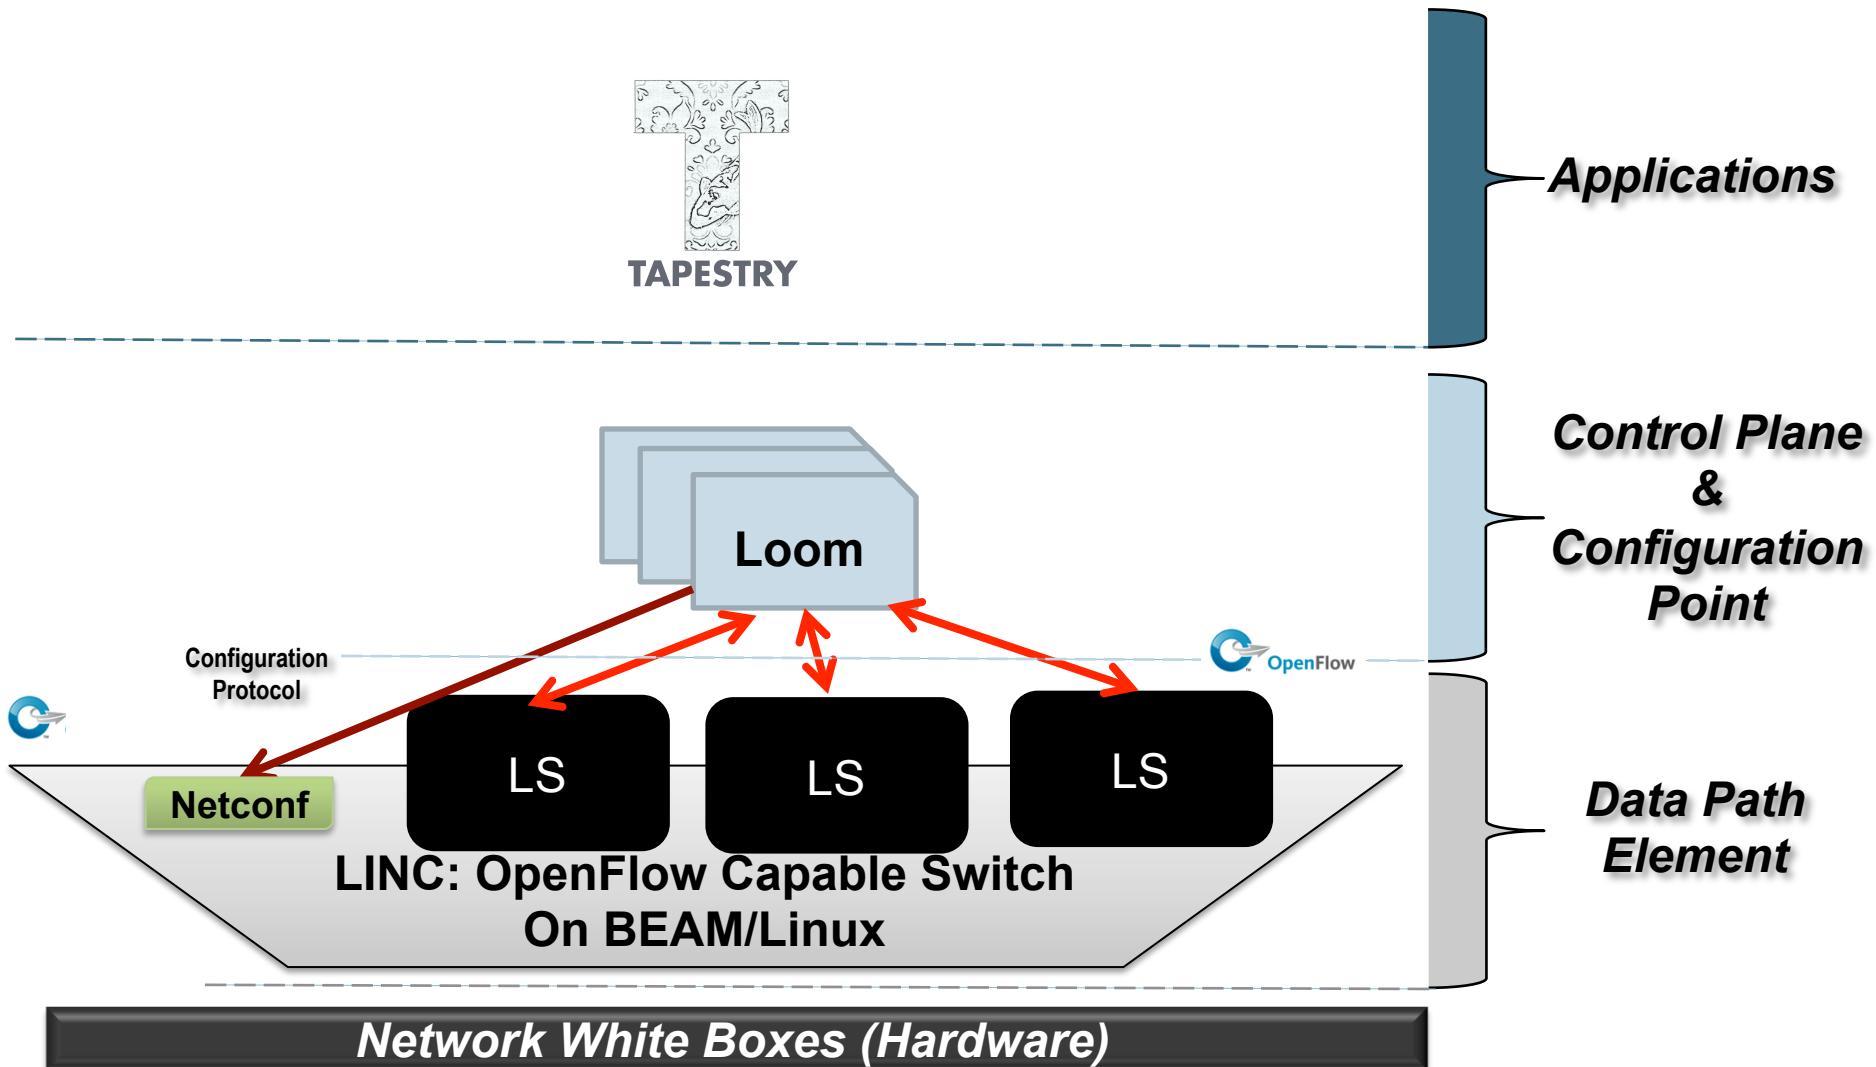
Loom Control Plane
POC
(Oct. 31, 2013)

Tapestry Network
Complexity
Analyzer v0.1
(Oct. 31, 2013)

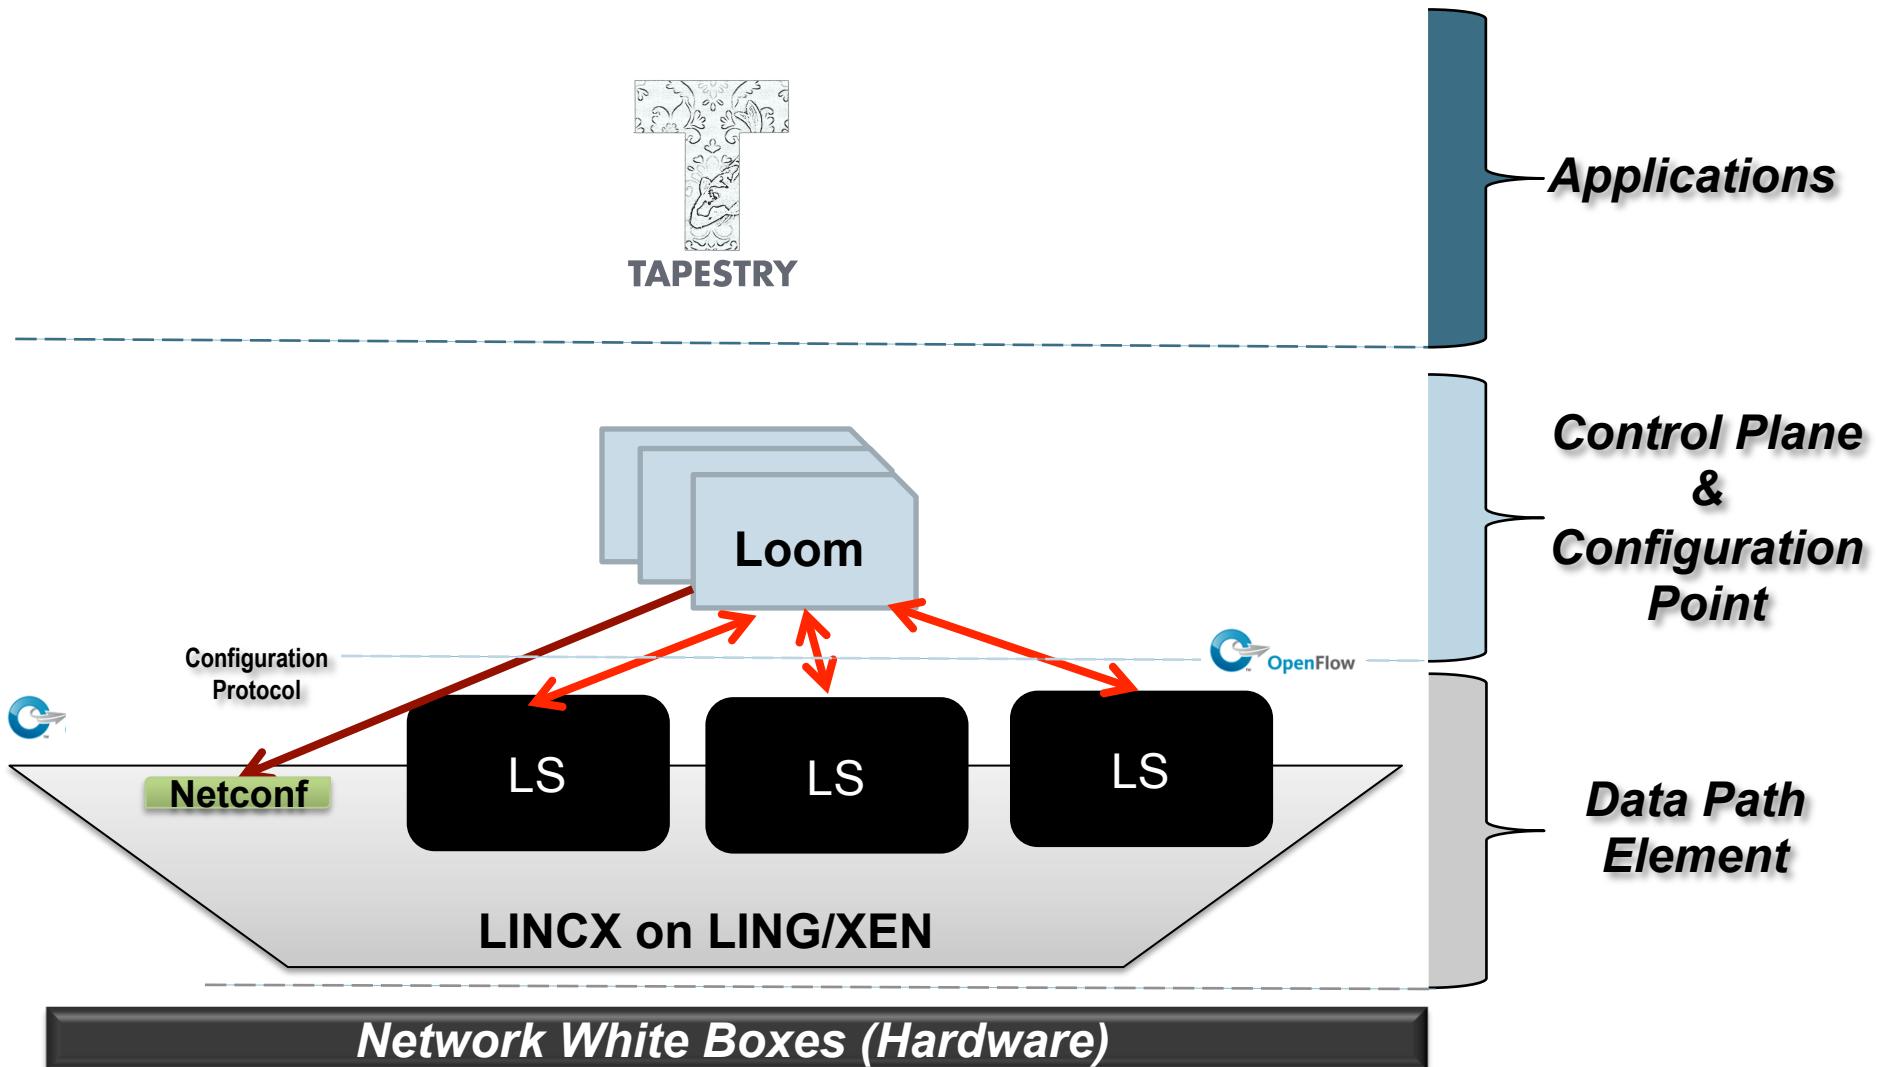
Loom v2
Scalable SDN
Control Plane
(Spring 2014)

Erlang
SOLUTIONS

Erlang = SDN + Big Data



Erlang = SDN + Big Data

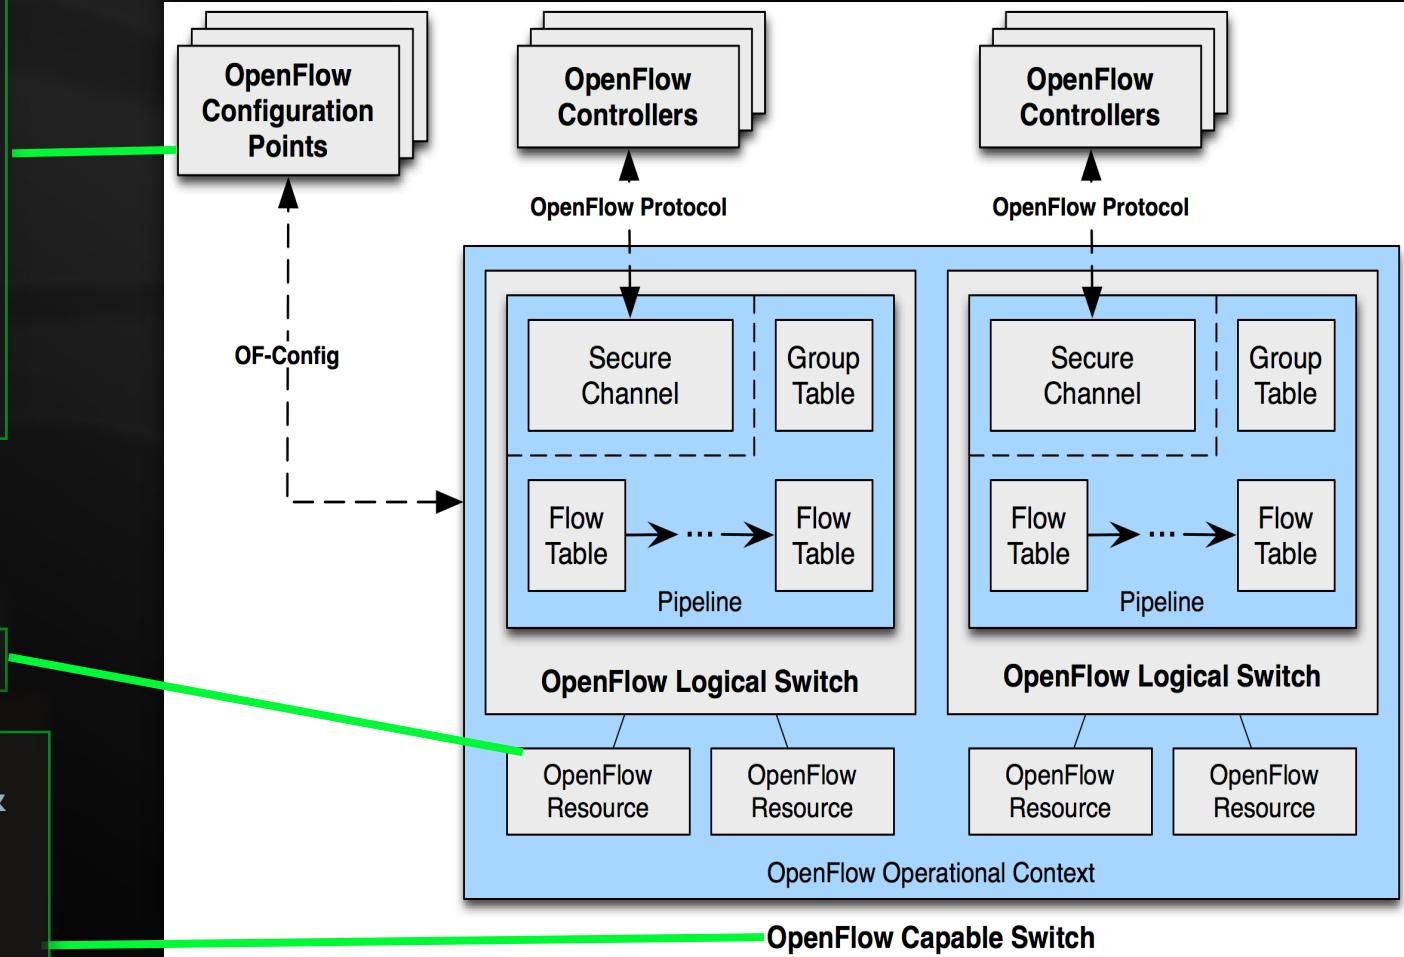


LINC Switch Architecture

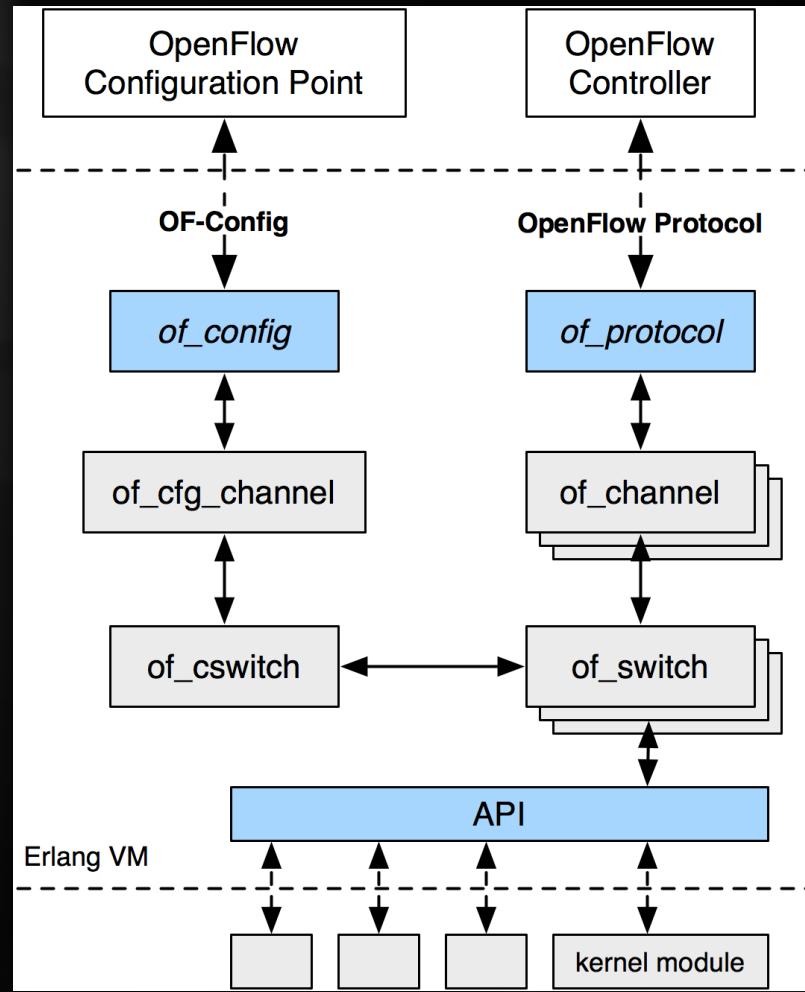
- Provisions OF Capable Switches
- Defines instances of logical switches
- Dispatches resources between switches

- Physical Ports

- Consists of ports & forwarding engine
- Container for multiple logical switches



Erlang Implementation Architecture



LINCX

Current results

3/5/14

Maxim Kharchenko
Cloudozer LLP

Overview

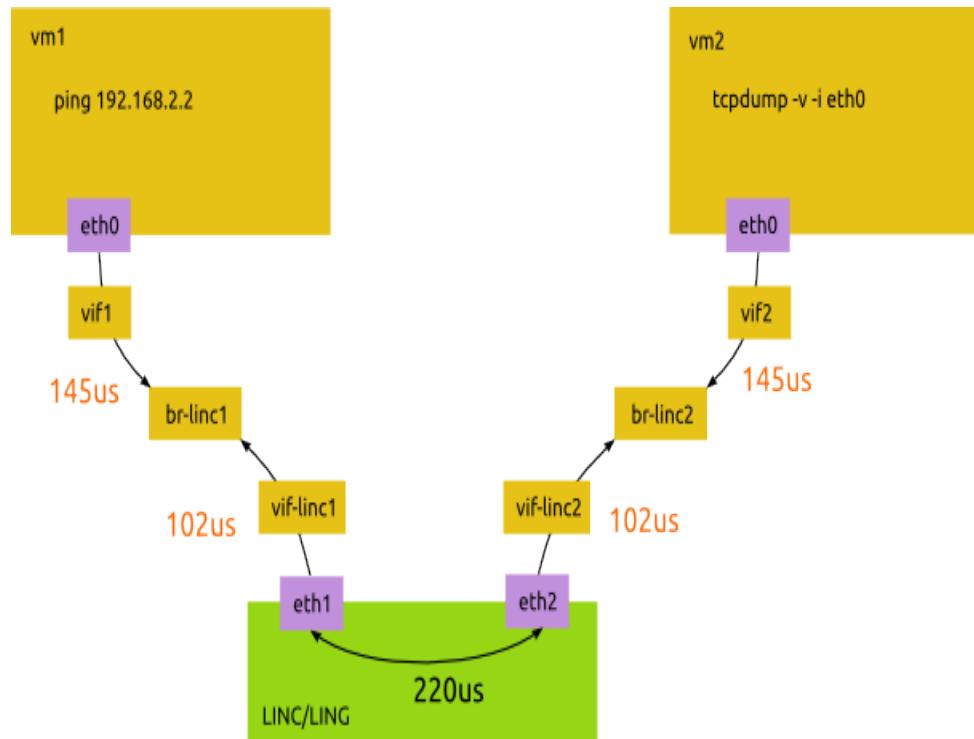
- Started December 9, 2013
- Initial scope = porting LINC to LING VM
- High degree of compatibility demonstrated for LING VM
- Extended scope = fix the LINC fast path
- Beta version of LINCX open-sourced March 3, 2014

Old backend (linc_us4)

- Trivial flows = 2 forwarding rules

	LINC/BEAM	LINC/LING
Max TCP throughput	140Mbit/s	200Mbit/s
Peak memory consumption	4.5GB	98MB
CPU utilization	200%	100%
Processing delay	-	220μs

LINC/linc_us4/LING latency

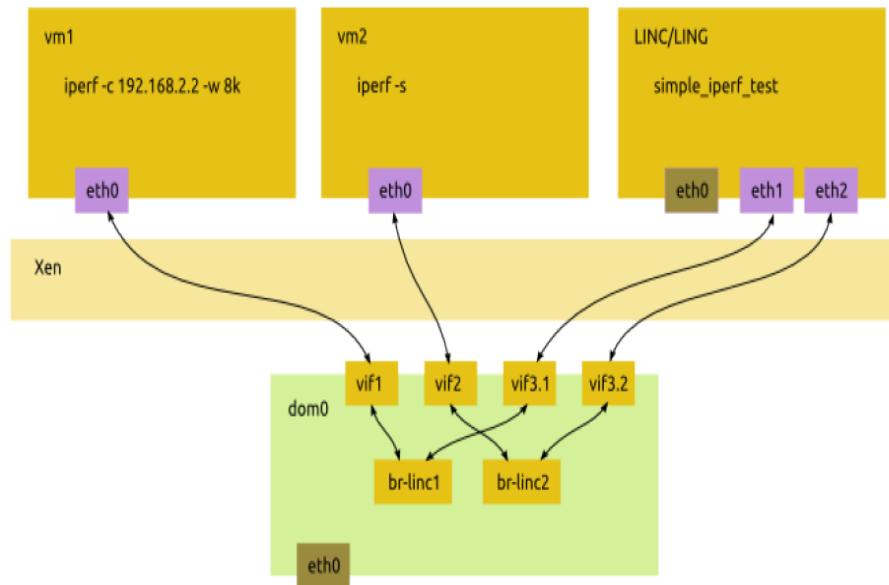


New backend (linc_max)

- Flow table = Erlang module with match() function
- Flow modules regenerated/compiled/reloaded for each flow_mod message
- GC avoided on the fast path
- Packet data (almost) never copied
- Memory footprint minimized

Live demo

- HP Proliant iLO 4 server
- Xen 4.2 hypervisor
- Bridged setup with two VMs:



linc_max 100x faster

- A non-trivial flow table with 256 rules:

LINC/linc_max/LING	
Max TCP throughput	2.5Gbit/s
Peak memory consumption	22M
Processing delay	3μs

Future

- Multicore
 - A switch per port? Per two ports?
- Faster software bridges
 - Replace generic Linux bridges with faster alternative
- PCI passthrough
 - Eliminate Linux completely from the packet path

The State of the Art...is Not

Data Platforms

Data Intensive Applications
(e.g. Hadoop, Apache Spark)

Manual System Abstractions

OS (e.g. Linux), Abstract Machines (e.g. JVM), Languages (e.g. Ruby), Clustering Frameworks (e.g. OpenStack), Databases/FileSystems (e.g. MongoDB, HDFS), HDN++

Fundamental Abstractions

Ethernet Frame

Virtual Machine

Increasingly Ubiquitous Hardware

Network Processors
(Broadcomm, Qualcomm)

Microprocessors
(x86, ARM)

Cloud
Private Cloud
Embedded Systems

The right perspective...and perfect timing

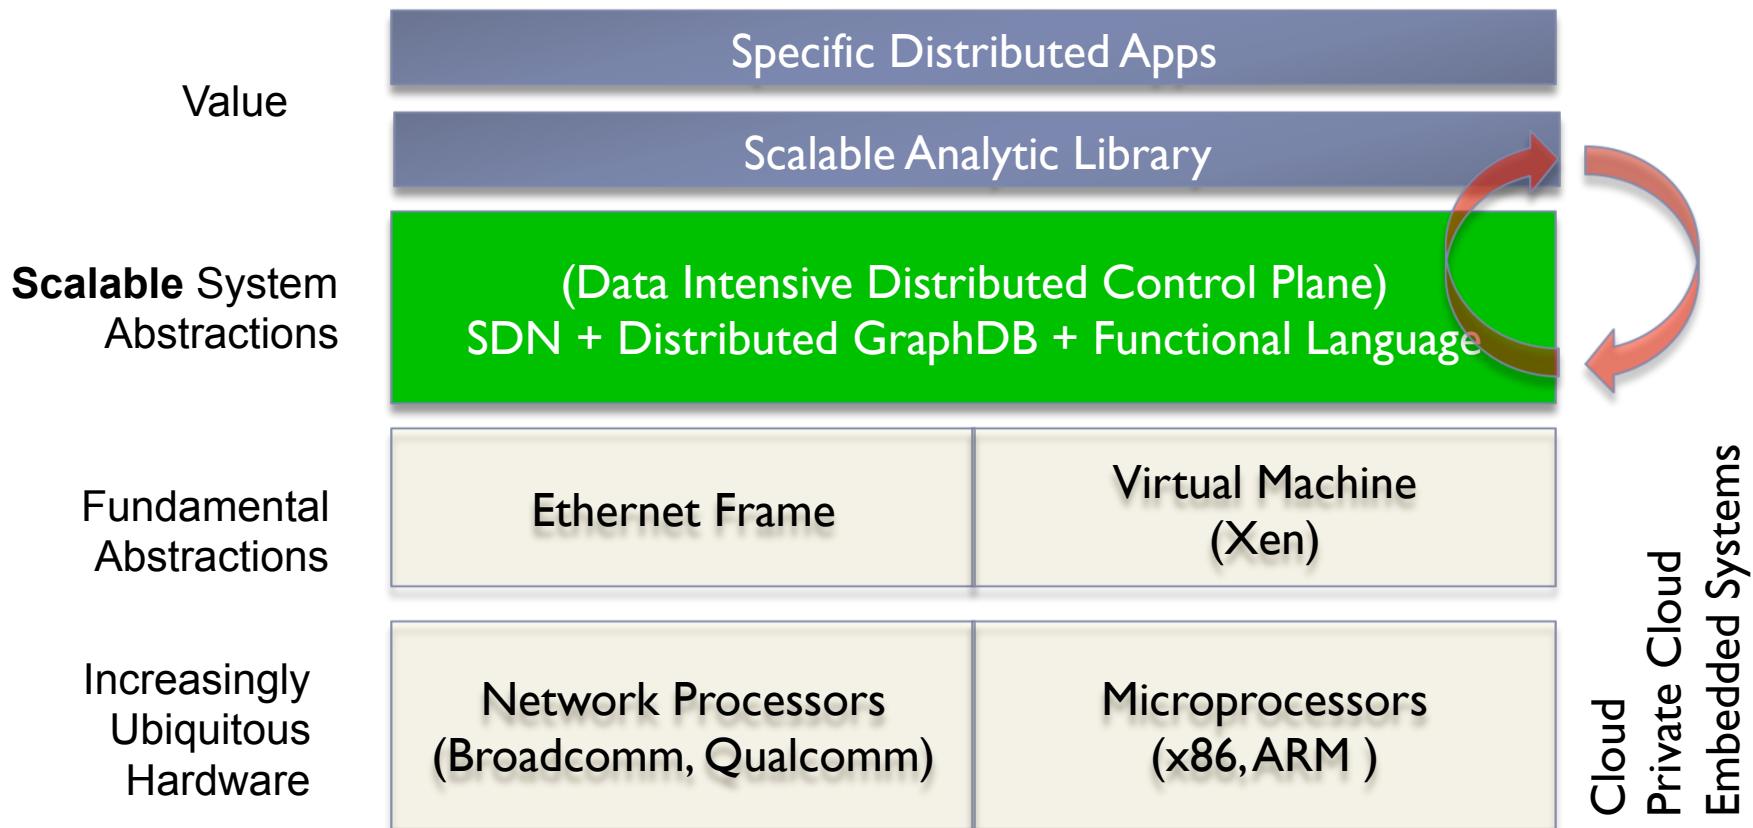


Ethernet + Multi-Core = A New Ubiquitous Machine

“The network is the computer” – John Gage (1984)



We need a better distributed system!



It is not a theory

Sample App

Tapestry
(Scalable Community Detection)

Getting there

Erlang/OTP + OpenFlow + Xen
(No Linux, No CloudStack, No Hadoop)

Fundamental Abstractions

Ethernet Frame

Virtual Machine

Increasingly Ubiquitous Hardware

Network Processors
(Broadcomm, Qualcomm)

Microprocessors
(x86, ARM)

Cloud
Private Cloud
Embedded Systems

Great but...

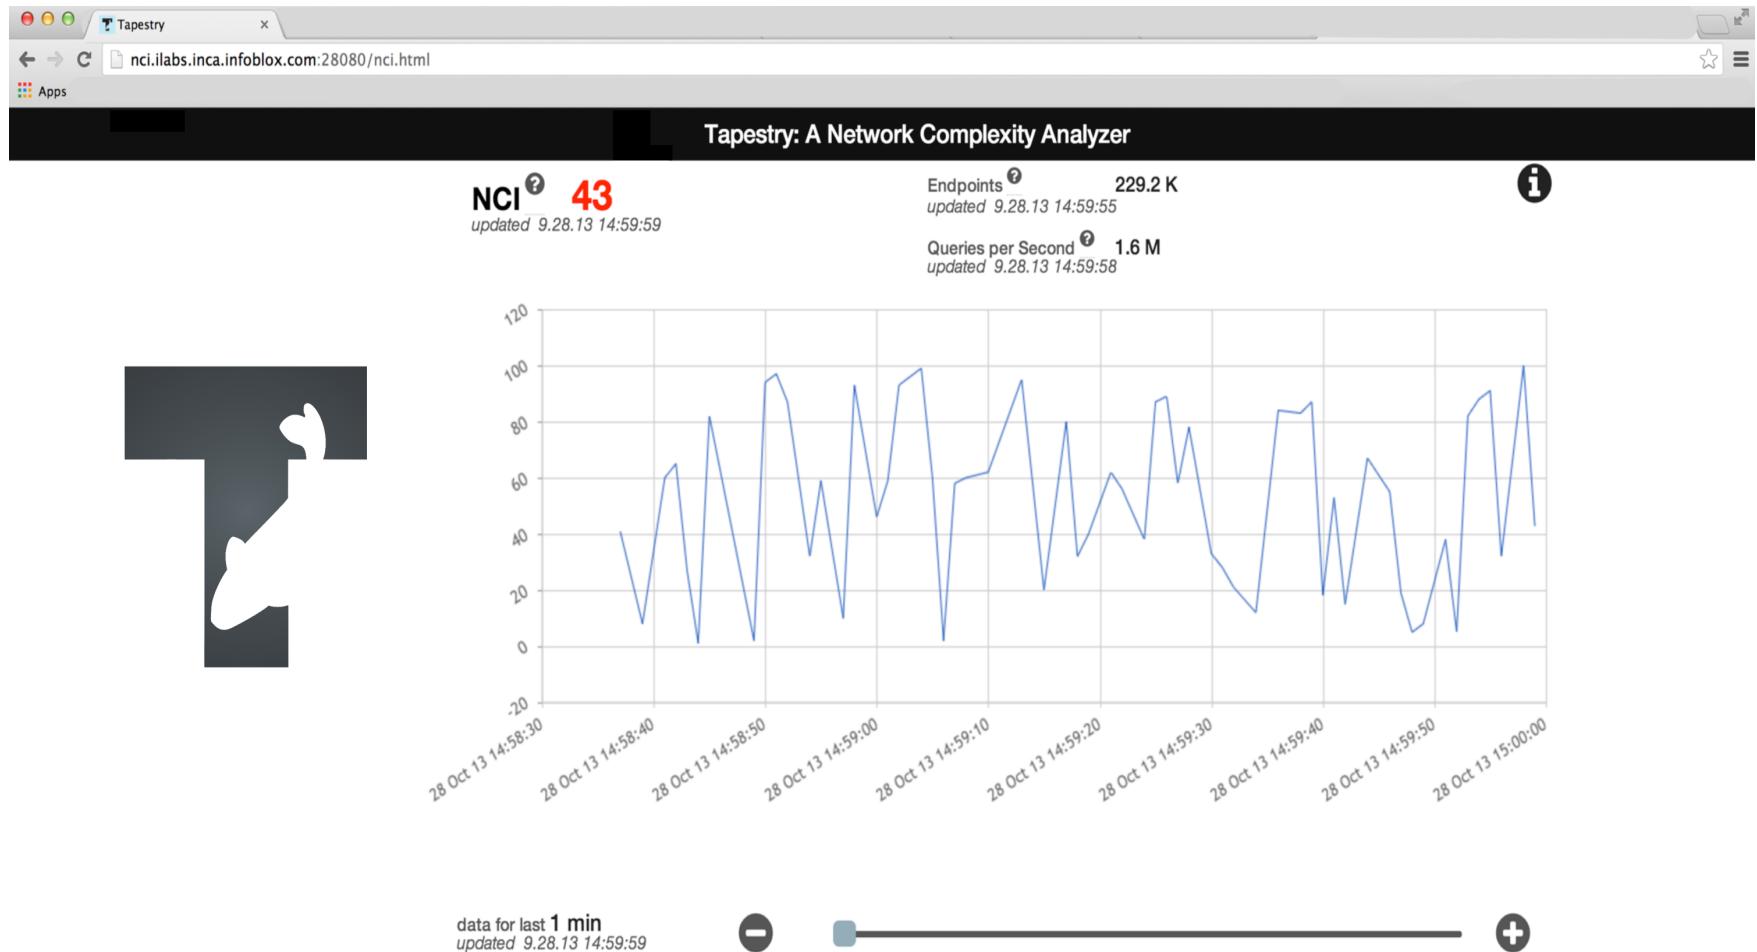


WhatsApp?

Is there
a **SINGLE NUMBER** that captures
NETWORK COMPLEXITY?

Easy to **COMPUTE**?
Easy to **UNDERSTAND**?





THE NETWORK COMPLEXITY INDEX

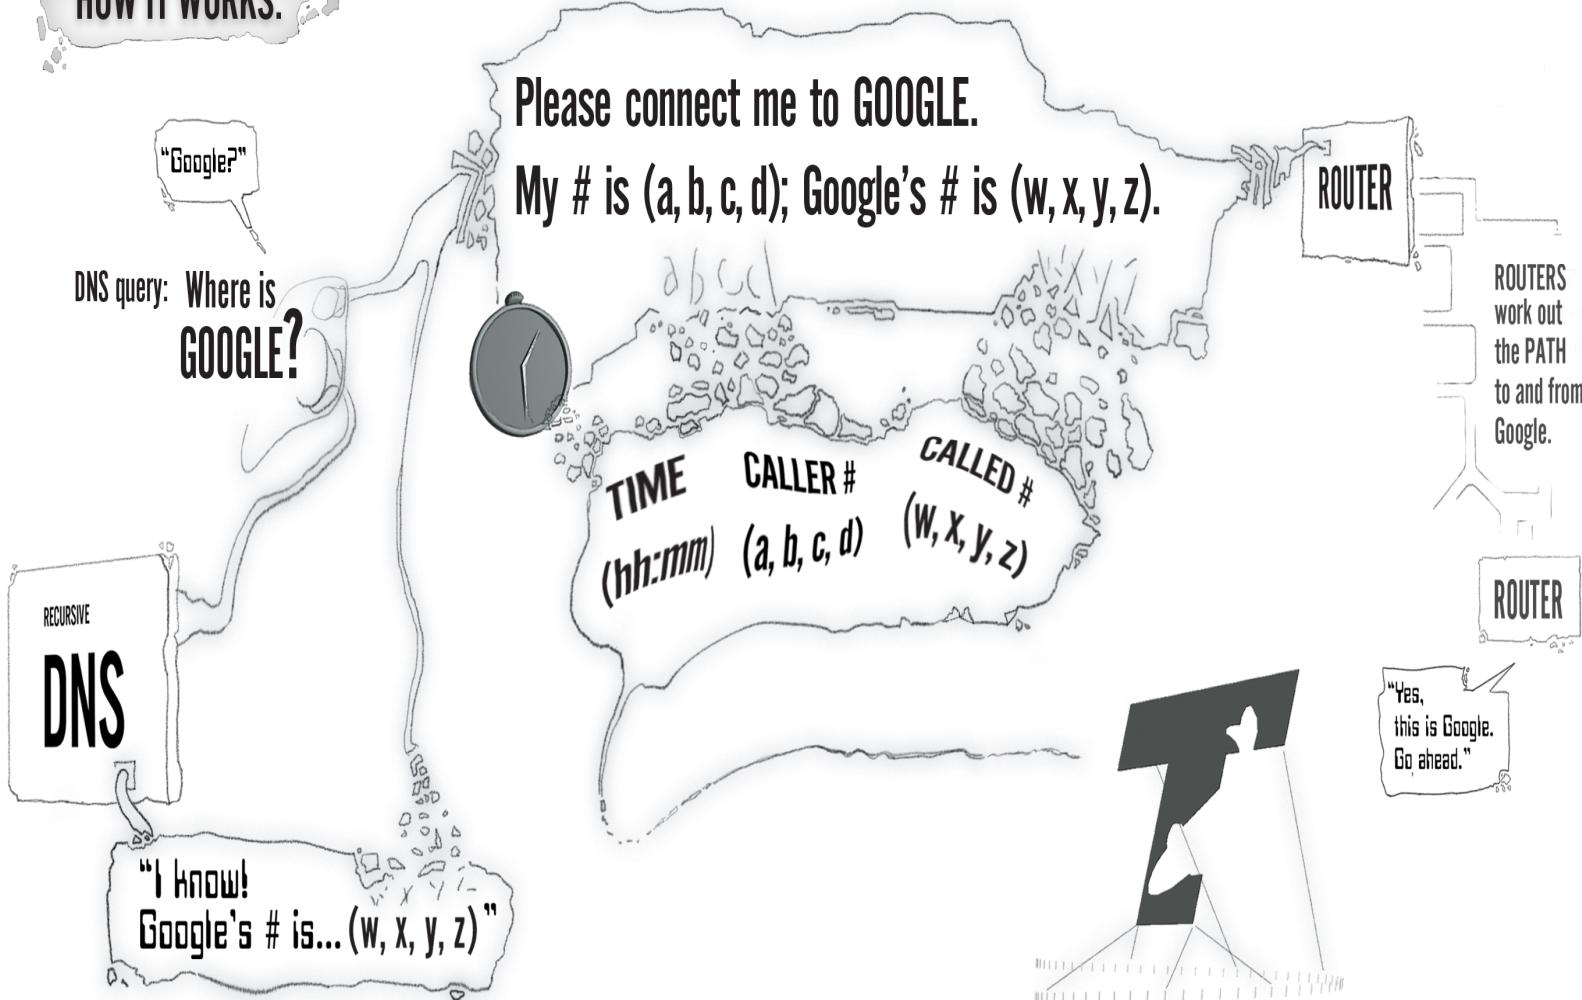


A measure
which rises with an increase in
the number of LARGE DISTINCT ACTIVITIES*
on a given network.

***LARGE DISTINCT ACTIVITY:**

Lots of participants
interacting toward a common goal.

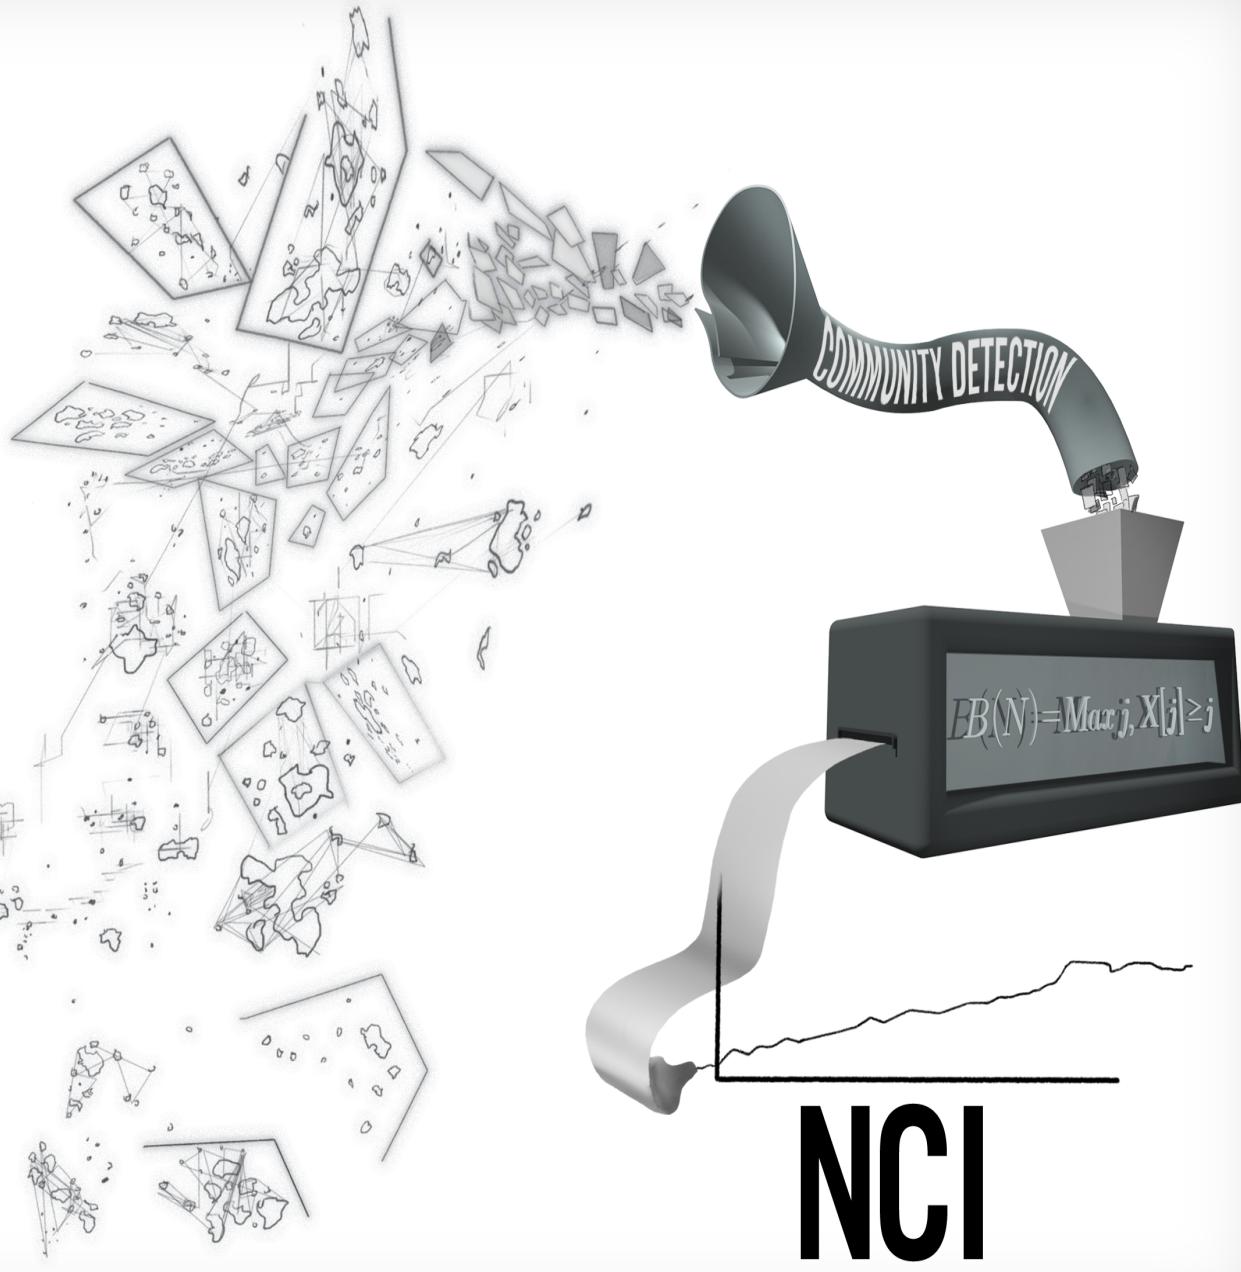
HOW IT WORKS.



TIME CALLER # CALLED #

(hh:mm) (a,b,c,d) (w,x,y,z)

(hh:mm)	(a,b,c,d)	(w,x,y,z)
(:)	(, ,)	(, ,)
(:)	(, ,)	(, ,)
(:)	(, ,)	(, ,)
(:)	(, ,)	(, ,)
:
:
:
:
:	''	'



The NETWORK COMPLEXITY INDEX

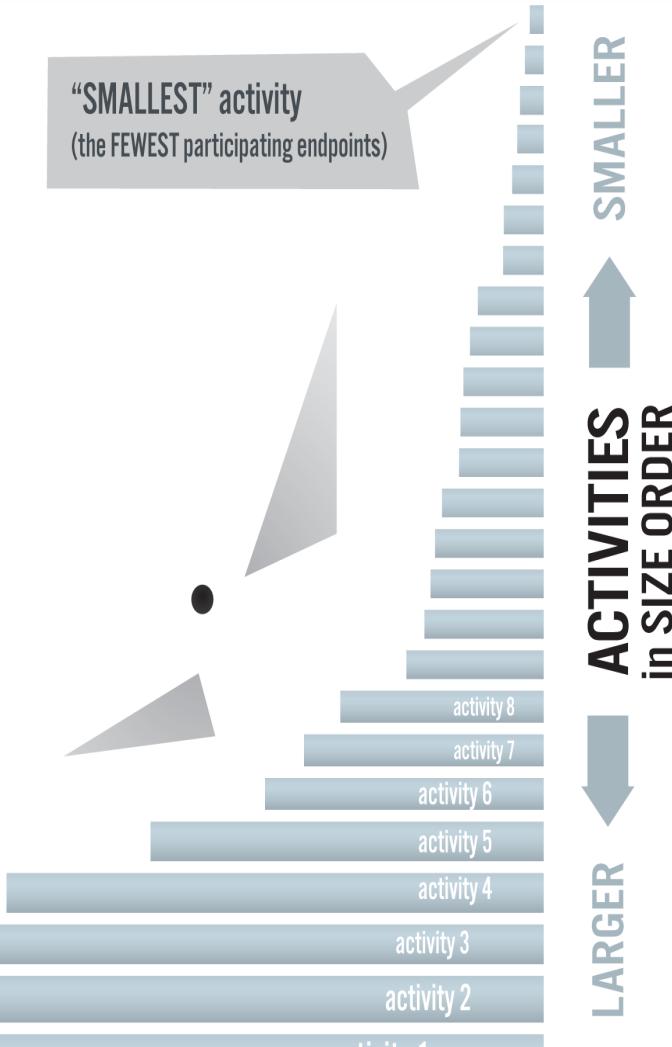
(NCI) The BALANCE POINT

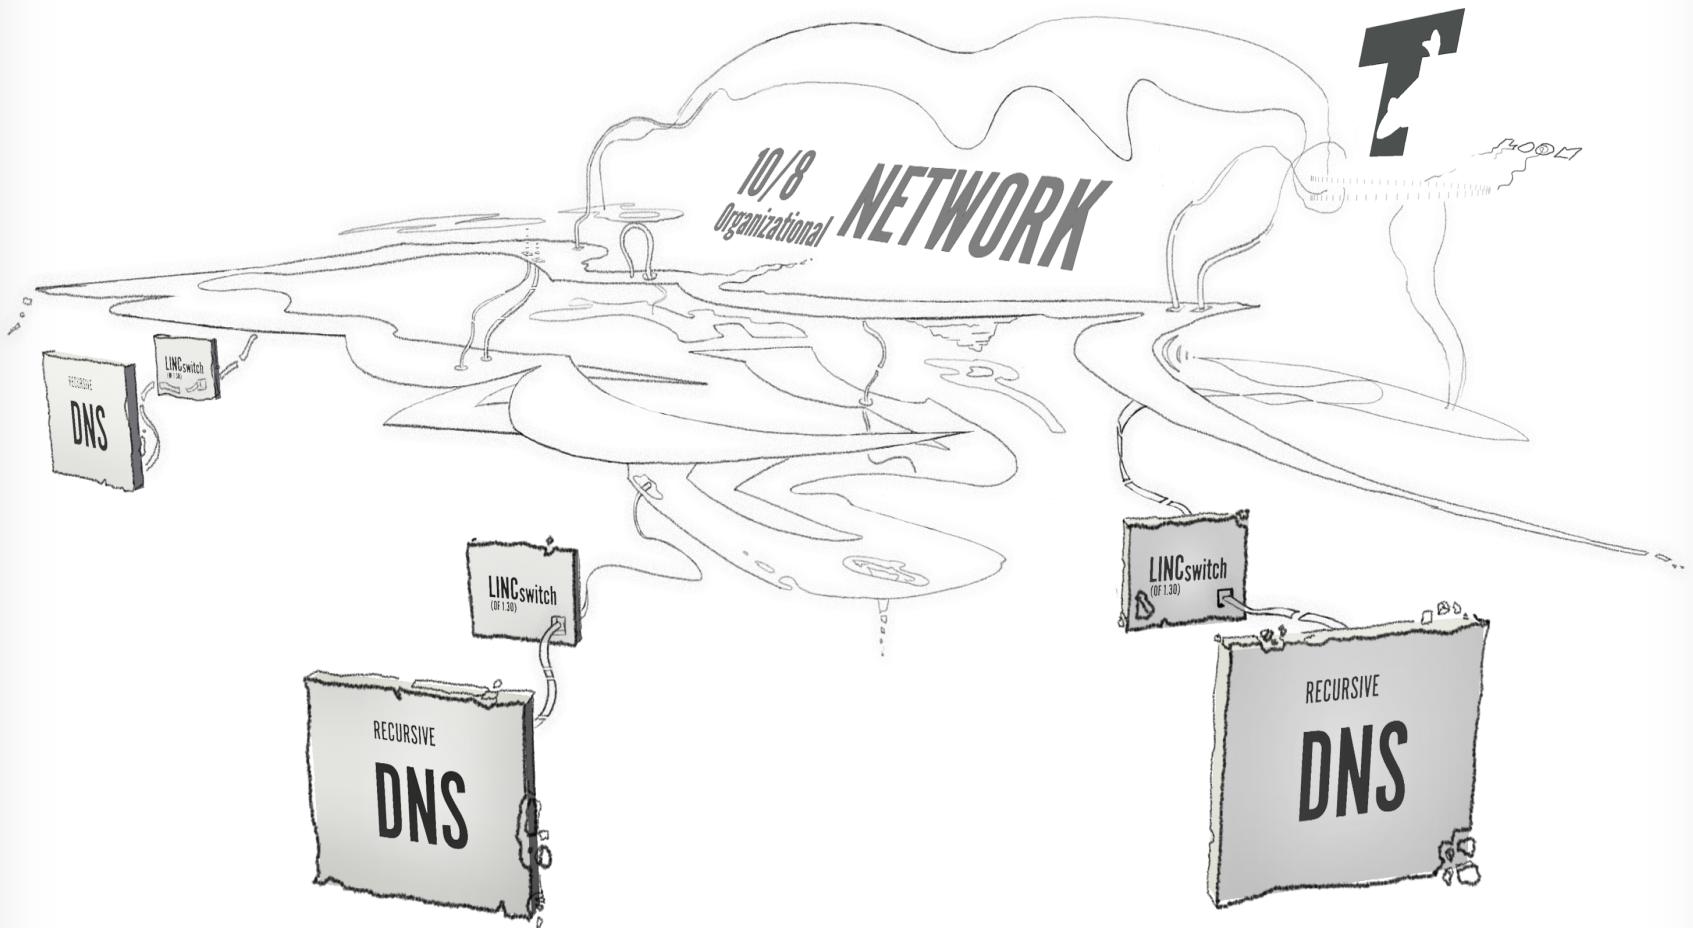
between
the SIZE of network activities
and
the NUMBER network activities.

“LARGEST” activity
(the most participating endpoints)

“SMALLEST” activity
(the FEWEST participating endpoints)

MANY ← # of ENDPOINTS → FEW







A New Machine?

IF



=



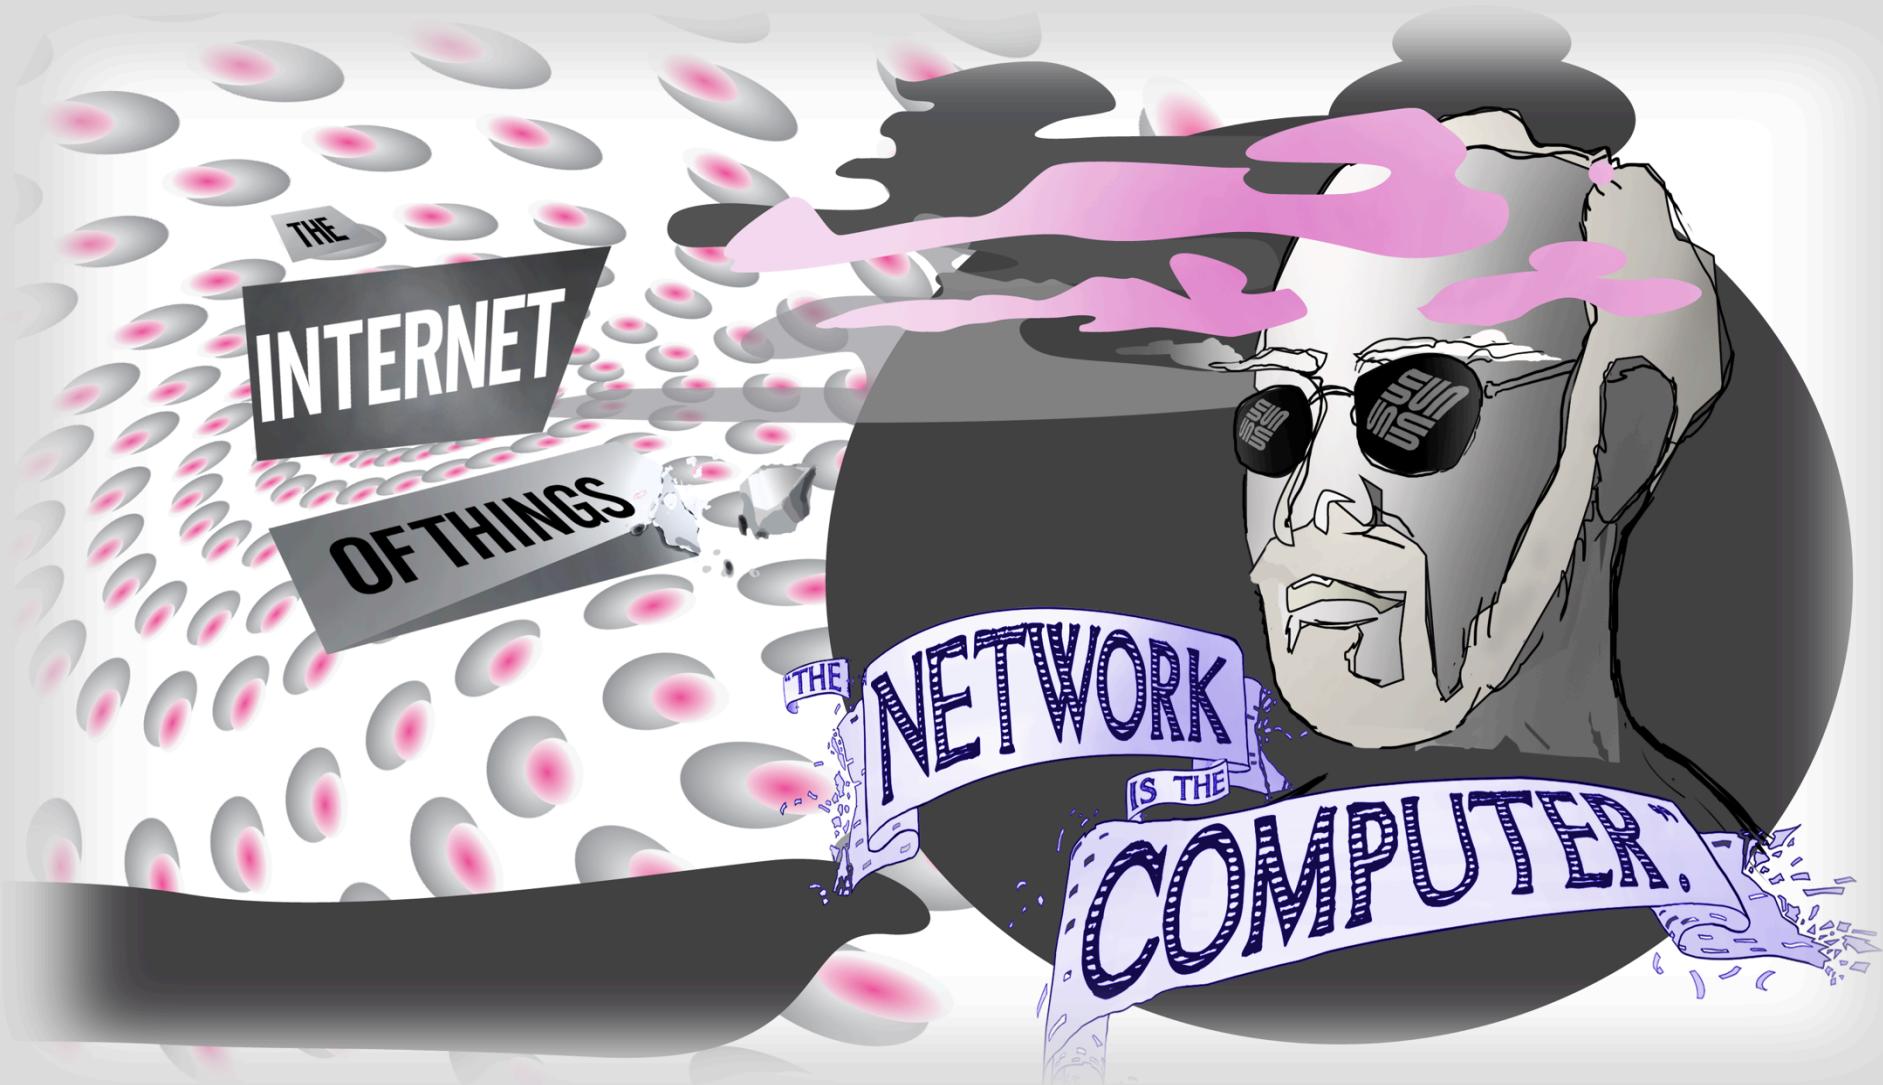
Then?



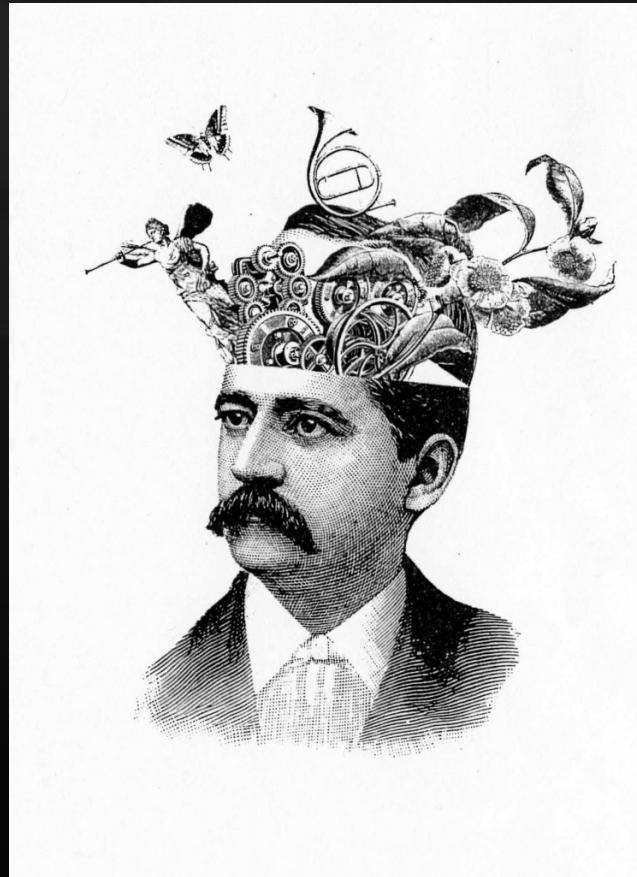
FLOW FORWARDING

© 2013 FlowForwarding.Org 36

We can look past The Cloud...



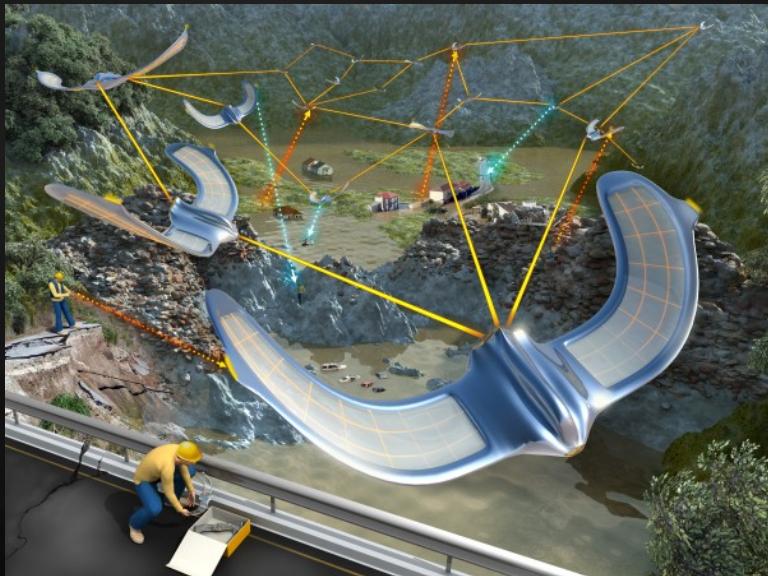
How can an Erlang/OTP developer program the entire network at runtime? (i.e. what abstractions, language extensions, mechanisms, etc.?)



For example, is I/O really a side effect?
Is computation king?



Or is communication just as fundamental?
Where are Joe's contract checkers??



Help us make it happen!

Email us @ info@FlowForwarding.org

<http://www.FlowForwarding.org>

