Collaborate on your eInfra, Only go Solo on Differentiators
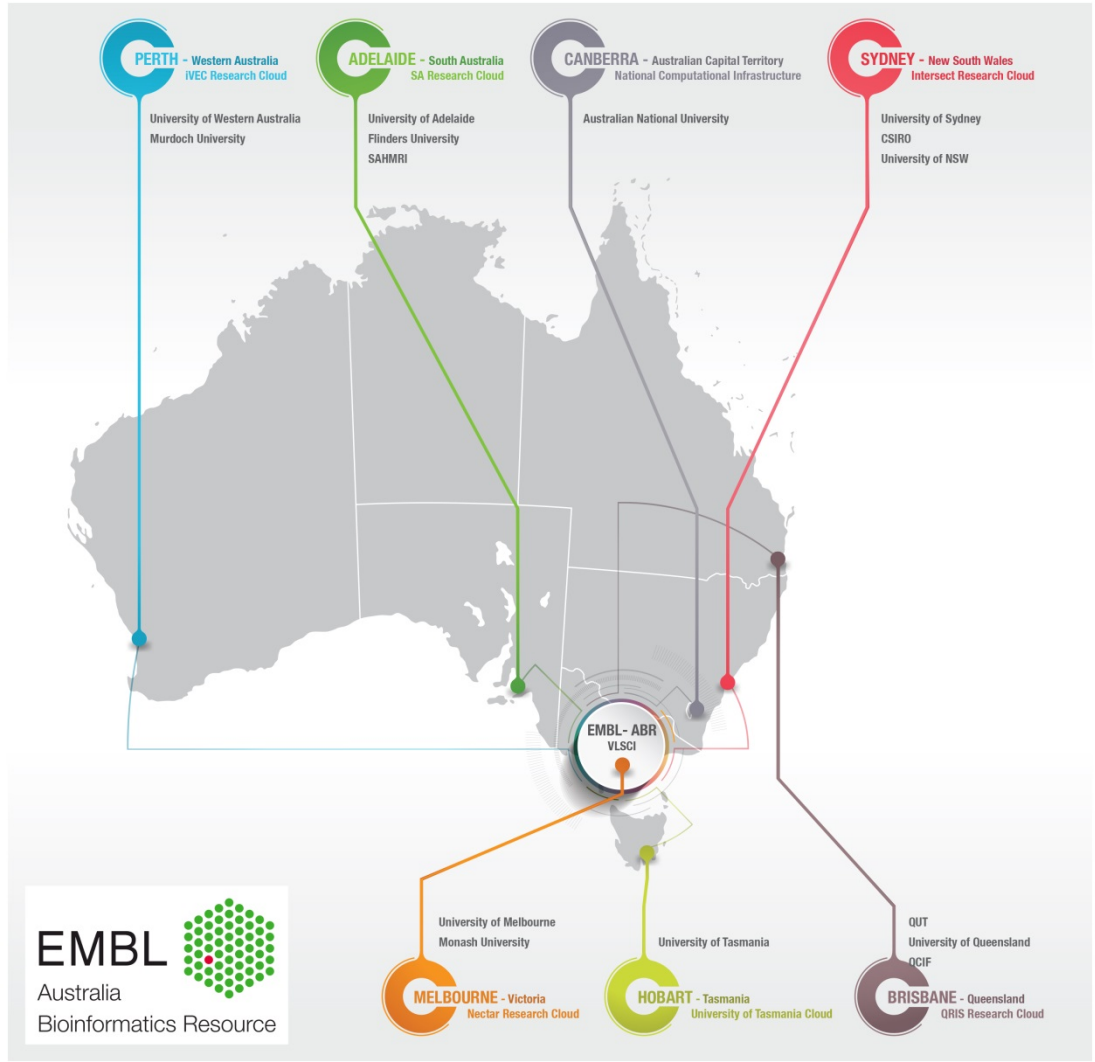
AboveTheNet workshop
October, 2016
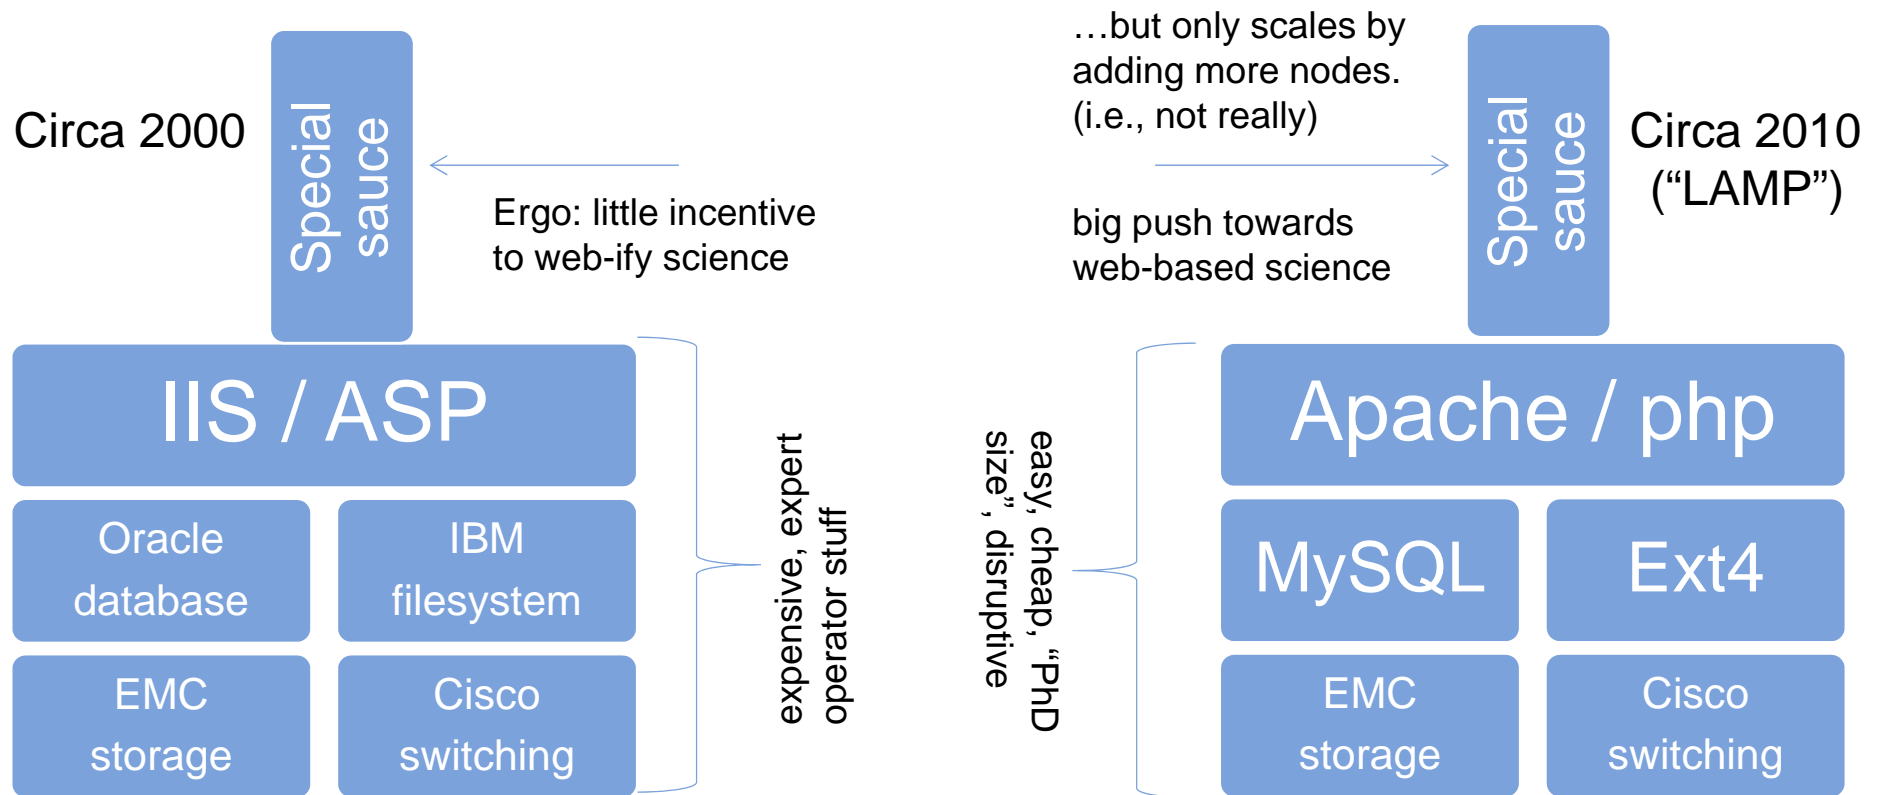
A little history.
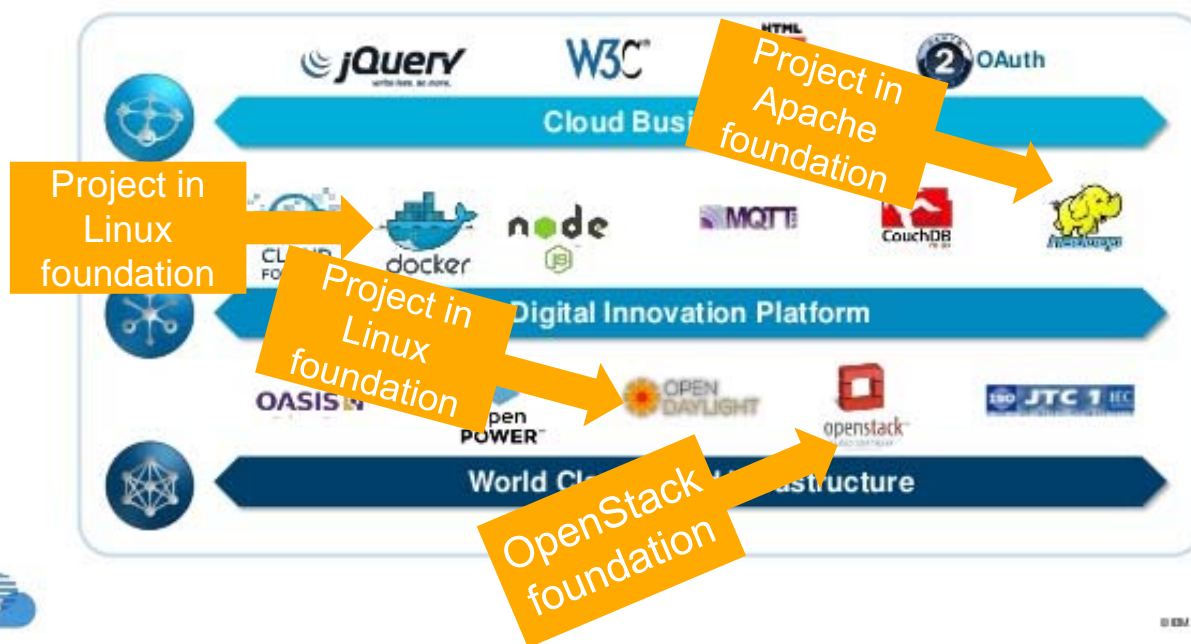## What sort of app delivery scaling has been available to science?

Circa 2000

Special sauce

← Ergo: little incentive to web-ify science

expensive, expert operator stuff

IIS / ASP

Oracle database	IBM filesystem
EMC storage	Cisco switching

…but only scales by adding more nodes. (i.e., not really)

big push towards web-based science →

Special sauce

Circa 2010 ("LAMP")

easy, cheap, "PhD size", disruptive

Apache / php

MySQL	Ext4
EMC storage	Cisco switching

Circa 2014 – a paradigm shift starts, bringing opportunity and challenge
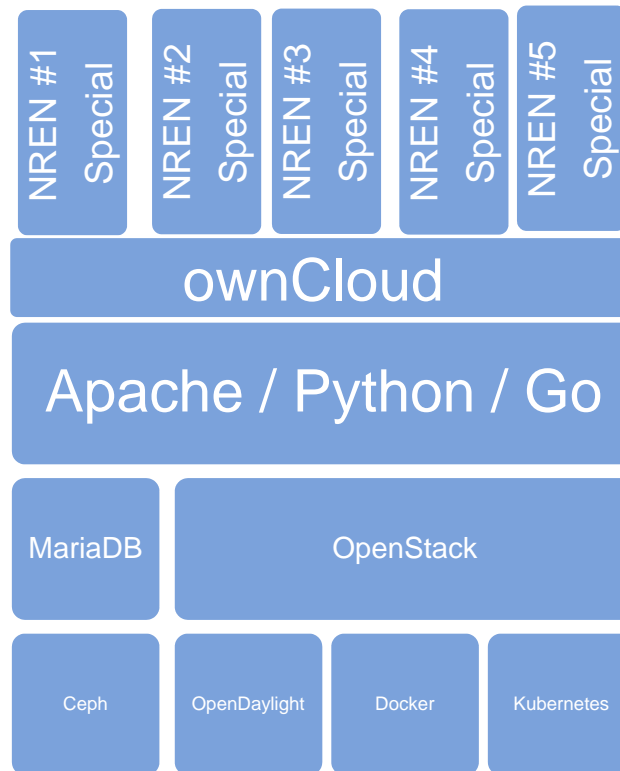

Suddenly a fully open source stack is available, that disrupts just about every scaling niche there was –

raw disk/Ceph for EMC

MariaDB/Galera for Oracle

OpenDaylight/Pica8 for Cisco

And it also scales cloudish, elastic and on-demand: OpenStack, Kubernetes, Docker, Dynamo. Cassandra

And sure enough, NRENs are aware. But something is funny...
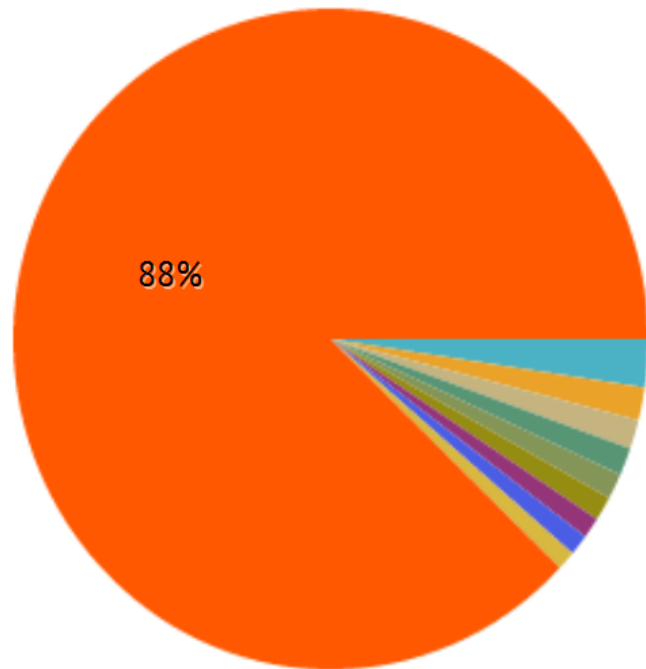
NREN response, Circa 2016

How big do you need to be? OpenStack example

Rank	Name	# Commits Newton	Revenue (M USD)
1	Red Hat	7745	1534
2	Mirantis	6082	10
3	HPE	3533	52110
4	Rackspace	3452	1794
5	IBM	2919	81740
	*independent	1941	
6	Intel	1567	55350
7	Fujitsu	1209	45440
8	NEC	1100	29520
9	Canonical	1077	66
10	SUSE	1006	200
11	Huawei	905	60839
12	EasyStack	857	4
13	ZTE Corporation	793	12095

Newton Rank	Name	# Commits Newton	Revenue (M USD)
14	99cloud	789	1
15	Cisco Systems	510	49240
16	Comcast	471	74510
17	VMware	400	6035
18	Nokia	351	25528
19	OpenStack Foundation	350	NONPROFIT
20	Universidade Federal de Campina Grande	320	ACADEMIC
21	AT&T	304	146800
22	Walmart	266	482100
23	Tesora	247	7
24	NTT	232	106700
25	Kylin Cloud	218	1
26	Midokura	203	5

Contribution by contributors

■	Andreas Jaeger
■	Emilien Macchi
■	Steve Martinelli
■	Alex Schultz
■	Jesse Pretorius
■	Paul Bourke
■	Yolanda Robla-Mota
■	Major Hayden
■	Stephen Finucane
■	others

How does OpenStack make decisions?

Great majority of technical decisions - made at technical level by the person contributing the code
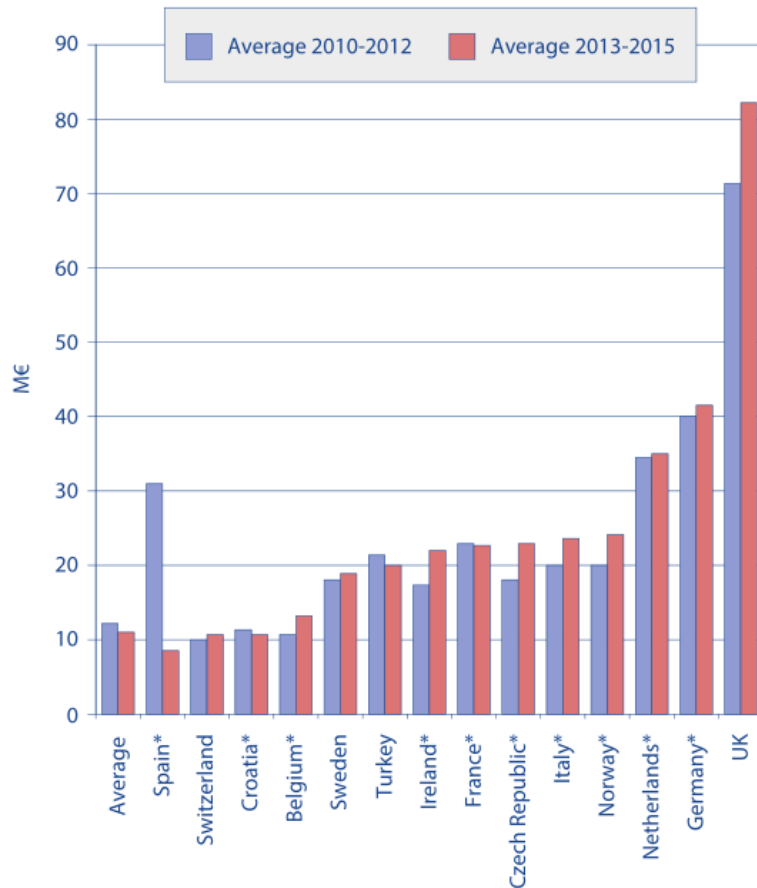
Tough decisions in a particular project - leader elected by the contributors of that project

Tough decisions across all/multiple projects - Technical Committee - elected from contributors to project

Representative body for user - User Committee - elected from users who do an activity from a define list

--committed *coders* matter more than *money*.

Graph 7.2.1 – Total budgets, 2010-2015 averages, GÉANT partner countries with annual budgets > 10 M€

How big are we as a collective? Financially:
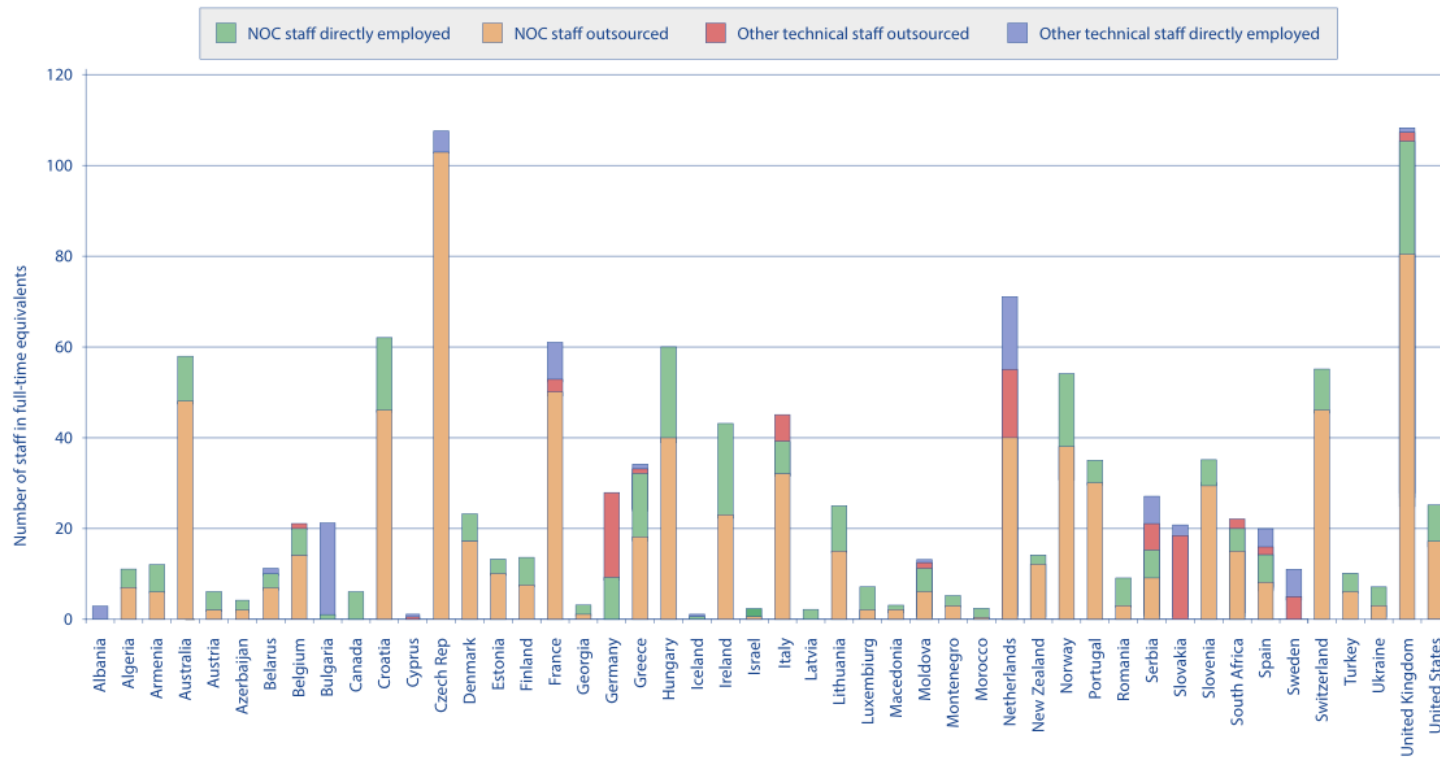
$$\sum \text{GEANT} \sim 370 \text{ M}$$

Add in Intenet2, RNP, AARNet
-- 500M€ conservative lower bound

How big are we as a collective? Coding manpower:

$$\sum compendium \sim 900 \text{ FTE techies}$$

Graph 7.1.2 – NREN technical staff in FTE

Legend: NOC staff directly employed · NOC staff outsourced · Other technical staff outsourced · Other technical staff directly employed

Y-axis: Number of staff in full-time equivalents

AARNet rule of thumb: 10% tech staff is coder

-- so ~90 coders all up

WE CAN DO THIS

...AND WE MUST DO IT.
COLLABORATION IS NO LONGER OPTIONAL

If we get it right, we'll be an excellent scaling partner for global science. If we get it wrong, science will drift to the cloud by osmosis. With no science left on our networks, irrelevance awaits.

ANY QUESTIONS?