



Michel Moriniaux

15/04/2016

Multivendor Datacenters

Disaggregating network services



Who's that dude?

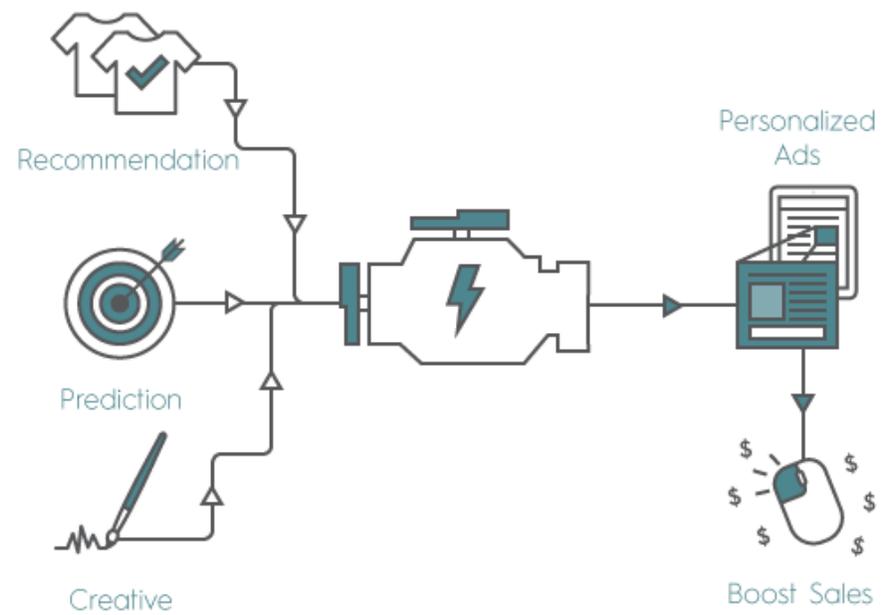
- 2 years at Criteo
- Network Infrastructure Team Lead

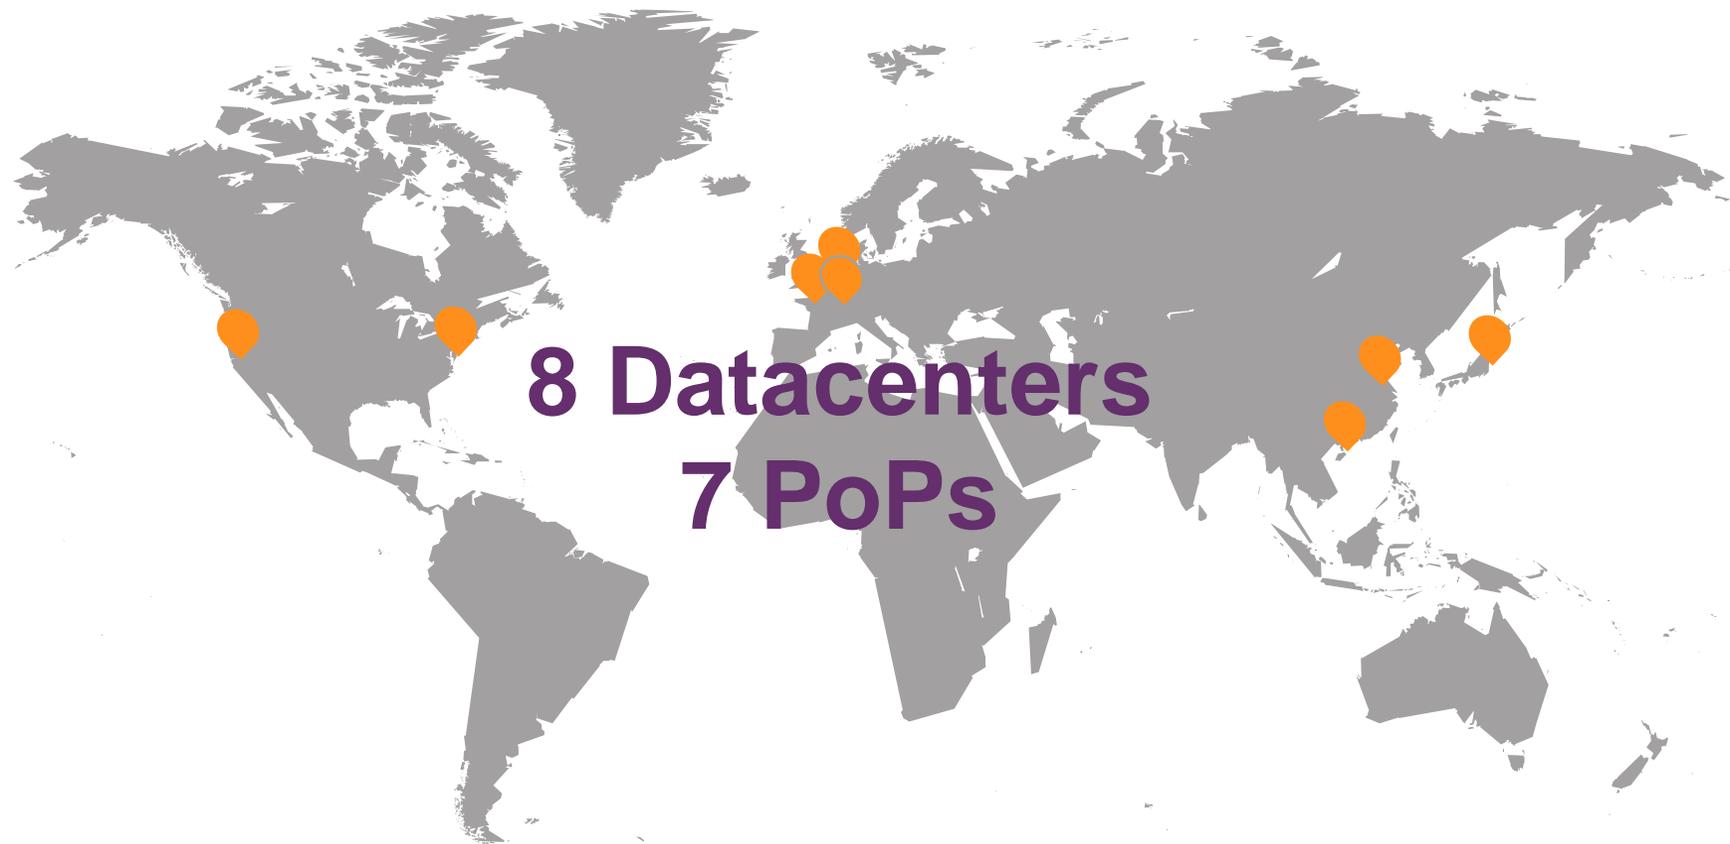
- 3 Missions:

- Network architecture & design
- Throwing out entrenched vendors
 - Automate myself out of a job

Transforming digital
advertising into a personal
experience.

About Criteo

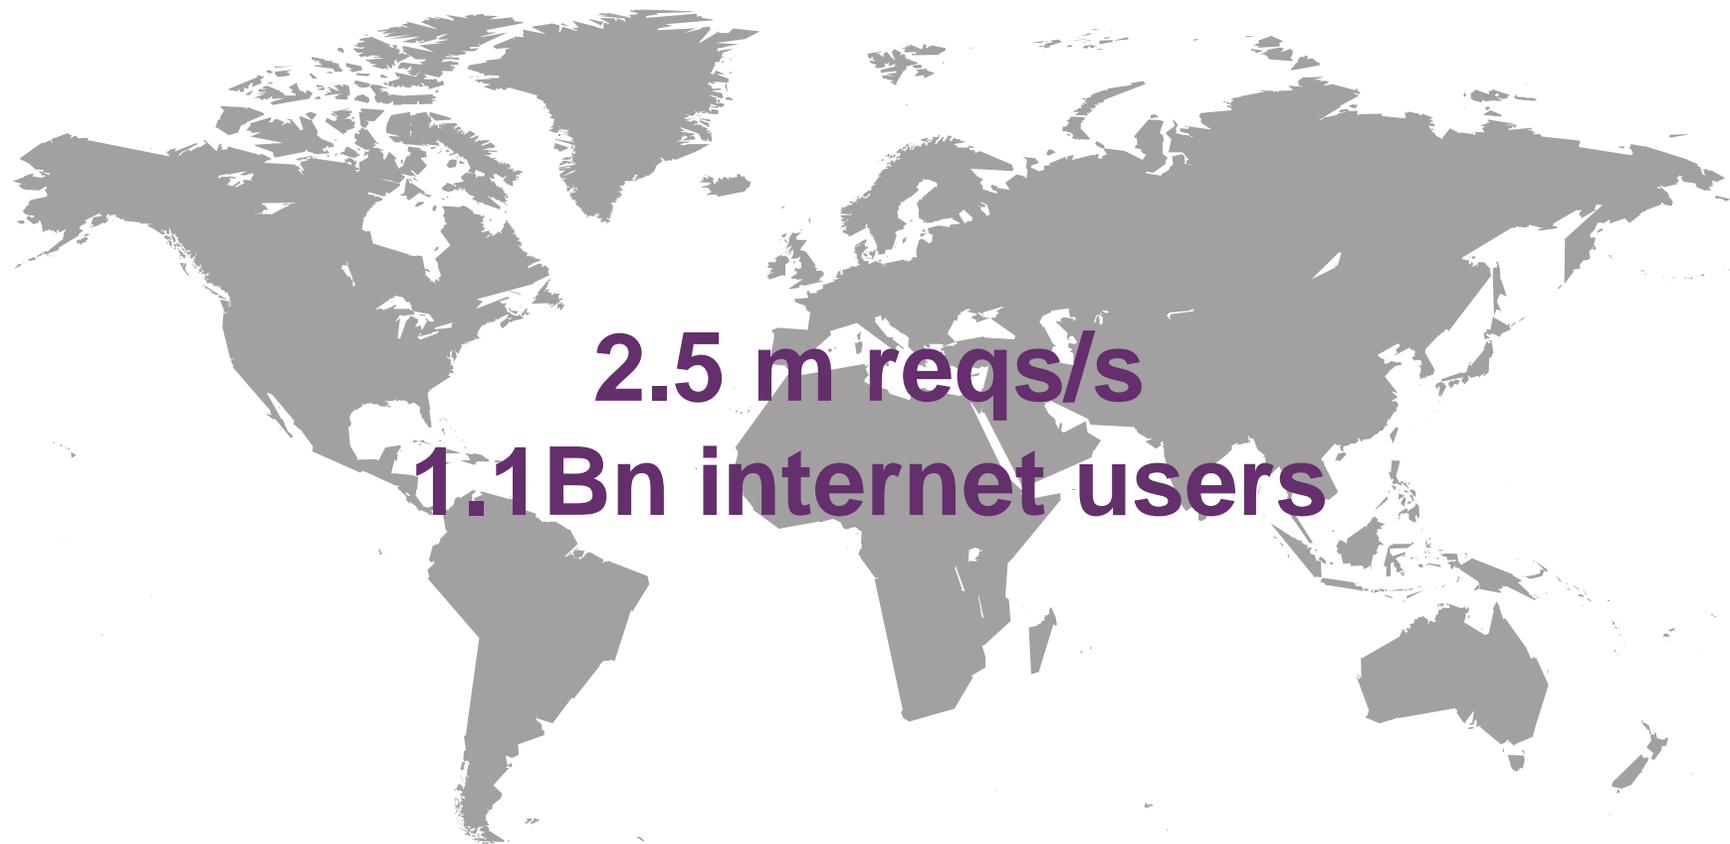












Multivendor in the DC



What usually happens when you have one vendor

- Build RFP
 - Consult
 - Choose

What usually happens when you have one vendor

Choose your dream Hardware

Send money to the Vendor

Sales moves to the Bahamas



Meanwhile...



- Please upgrade
- Known bug, goto 1
- Yeah, that's a feature

How do you solve this problem

- Story time! How we built PA4
 - Criteo's 2nd HPC cluster
- 40G non-blocking L3 Clos fabric

How do you solve this problem

- Design:
 - Proprietary features do not save lives
- Find the smallest set of features possible
 - IEEE and IETF are your friends

How do you solve this problem

- L3 Clos Fabrics
- One feature: BGP
- All the rest is useless

How do you solve this problem

- During the RFP process try to qualify at least 2 vendors
 - People will say you're crazy:
 - Twice the work, twice the problems

How do you solve this problem

- No vendor is used to this
 - Most will work with you
- Some will spread rumours of unfairness

How do you solve this problem

- Be Transparent
- Brief your vendors on the qualification results
 - Share anonymized scores

How do you solve this problem

Scoring table POD

feature				
BGP	1	1	2	-1 (green)
LLDP	1	1	1	1
IPV4	1	1	1	1
IPv6	2	1	2	-2
SPAN	2 (rspan)	1 (simple)	2 (rspan)	2 (rspan)
Flow	2 (Sflow)	2 (sflow)	2 (Jflow)	2 (Sflow)
L2 overlay	2 (VxLan)	2 (VxLan)	2 (VxLan)	-2
API	2 (API -shell)	1 (API)	2 (XML-RPC -shell)	1 (XML API)
CLI	2	2	2	1 (hybrid)
Code Stability	0	0	0	0
Code maturity	1	1	2	-1
...				
Auto Install	2 (ZTP)	2 (ZTP)	2 (ZTP)	-1 (dhcp)
Debugs	1	1	2	1
Documentation	1	-1	2	-1
User Base	1	1	2	-1
User Community	1	-1	2	-1
CPU	2 (x86)	0	2 (x86)	-1 (PPC)
redundant PS	2	2	2	2
1U	2	2	2	2
port capacity	2	2	2	2
Strong points	Deep Buffers automation	support	code maturity features	Price HPC experience
Weak points	Price	Documentation userbase only china		BGP code 09/2014
Total	30	21	36	4

Scores:

score	meaning
-2	unsupported
-1	low confidence
0	unknown
1	confidence
2	high confidence

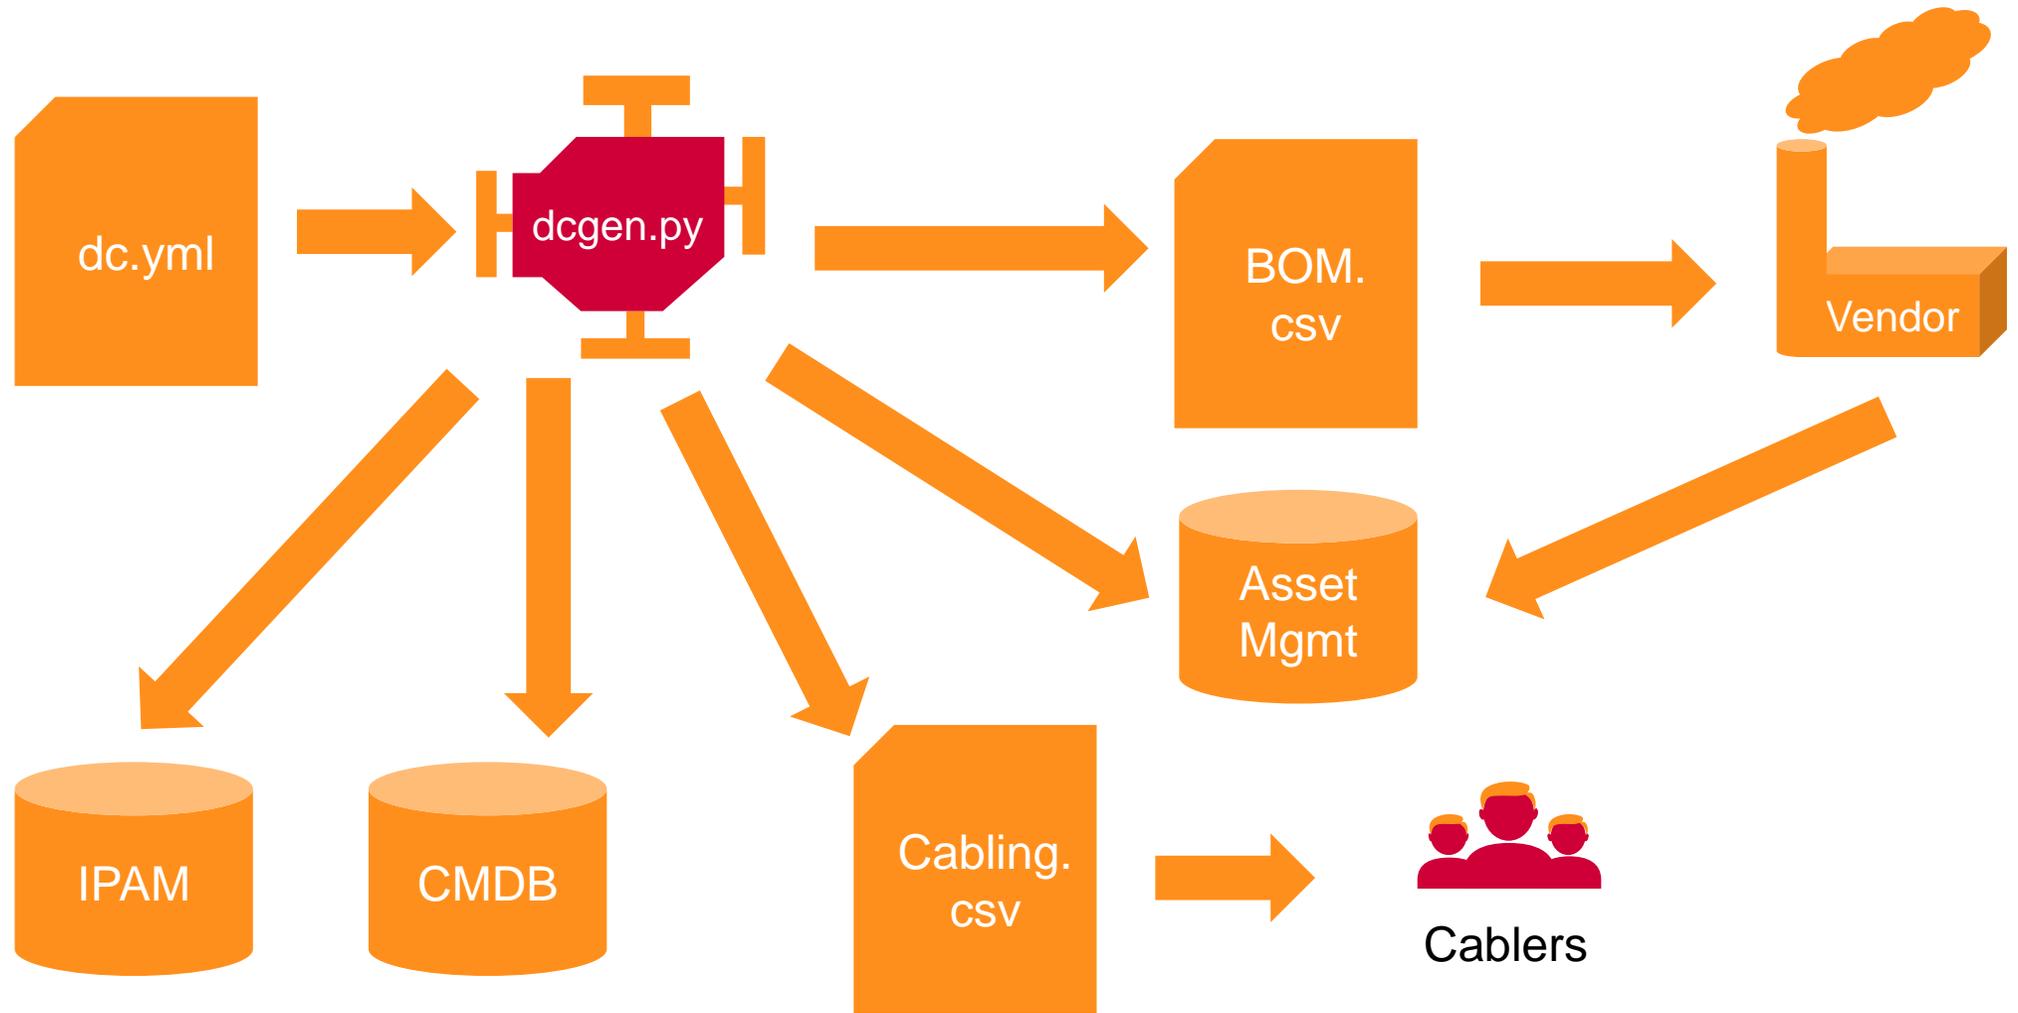
How do you solve this problem

- Chose the 2 best that fit your requirements
 - When you have time qualify more
- Give the power back to your purchasing department

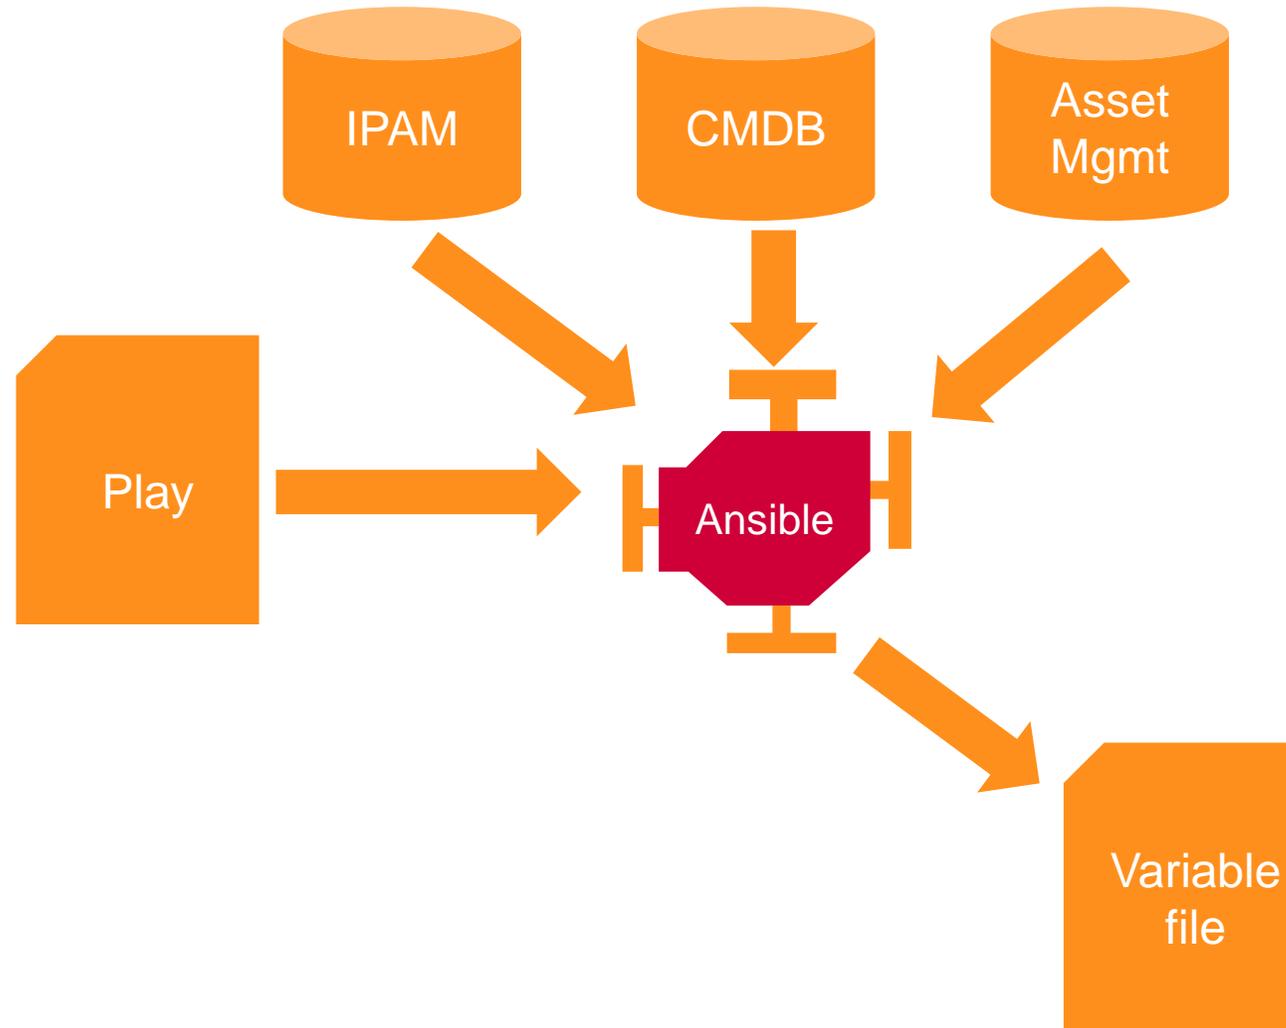
- The only way this can work is through automation
 - If you use the CLI you will never scale
- CLI is the last resort when stuff is broken

- Find an abstraction layer that fits you
 - Ansible var files
- Build a DC factory

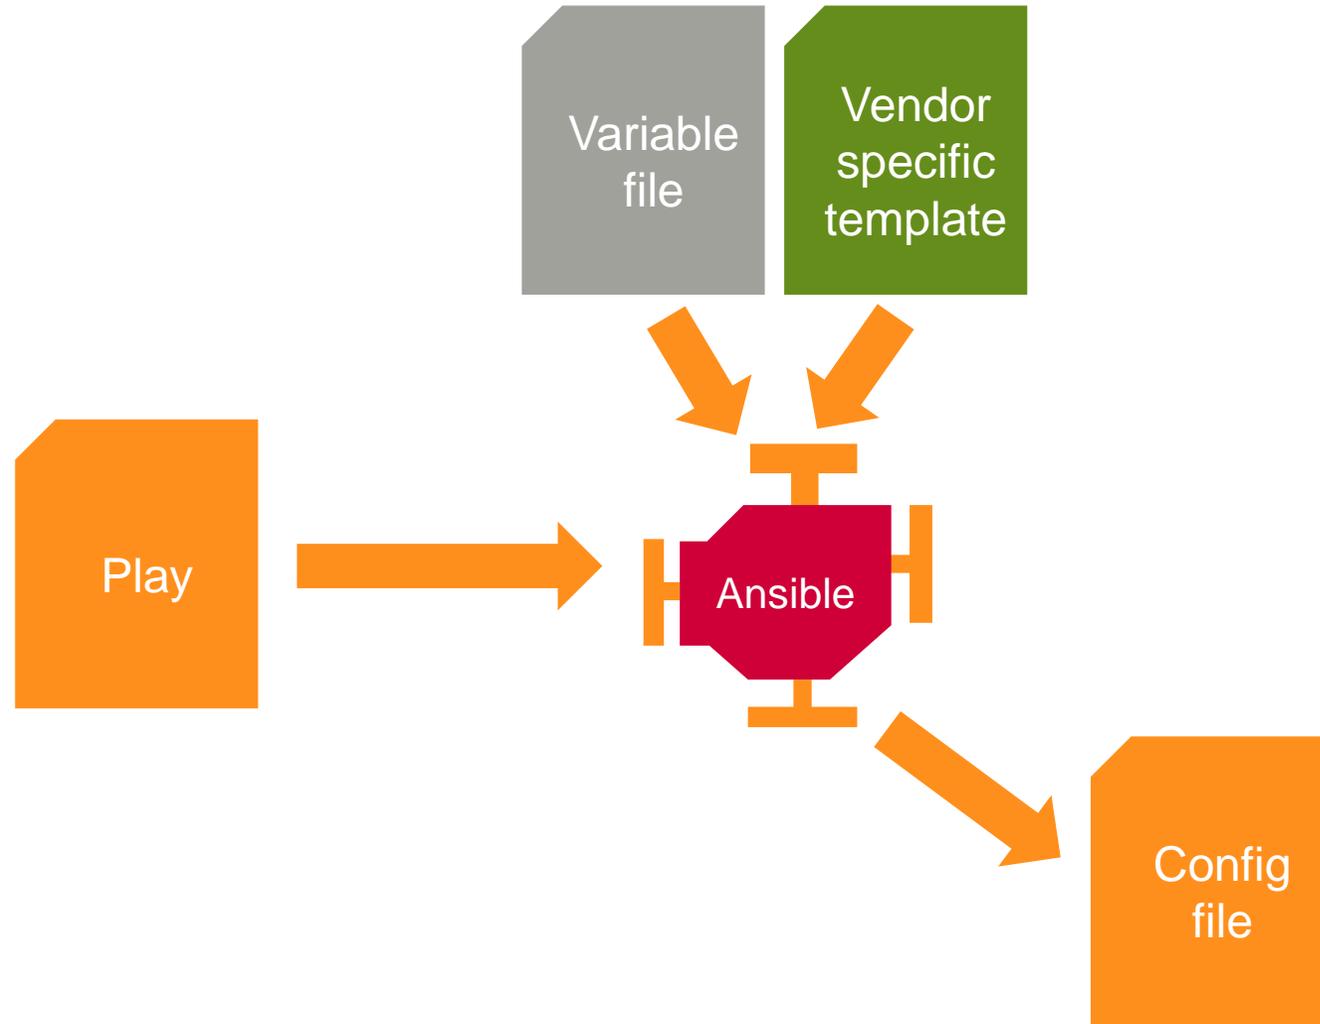
Automation is key



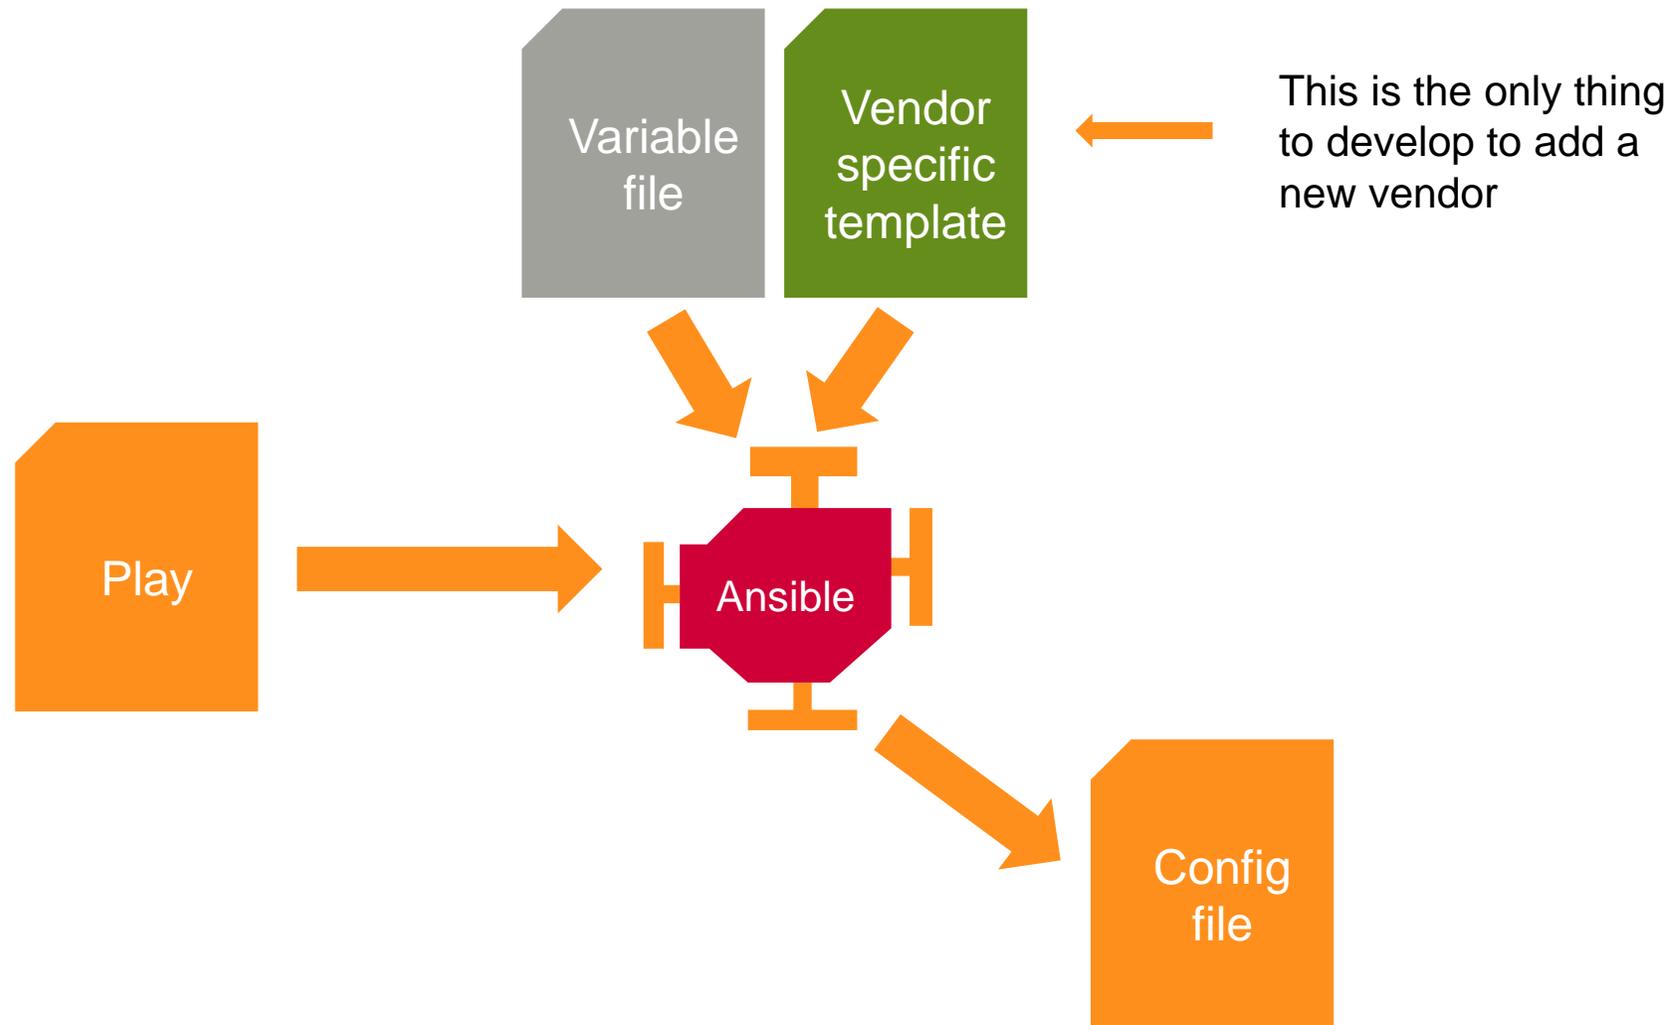
Automation is key



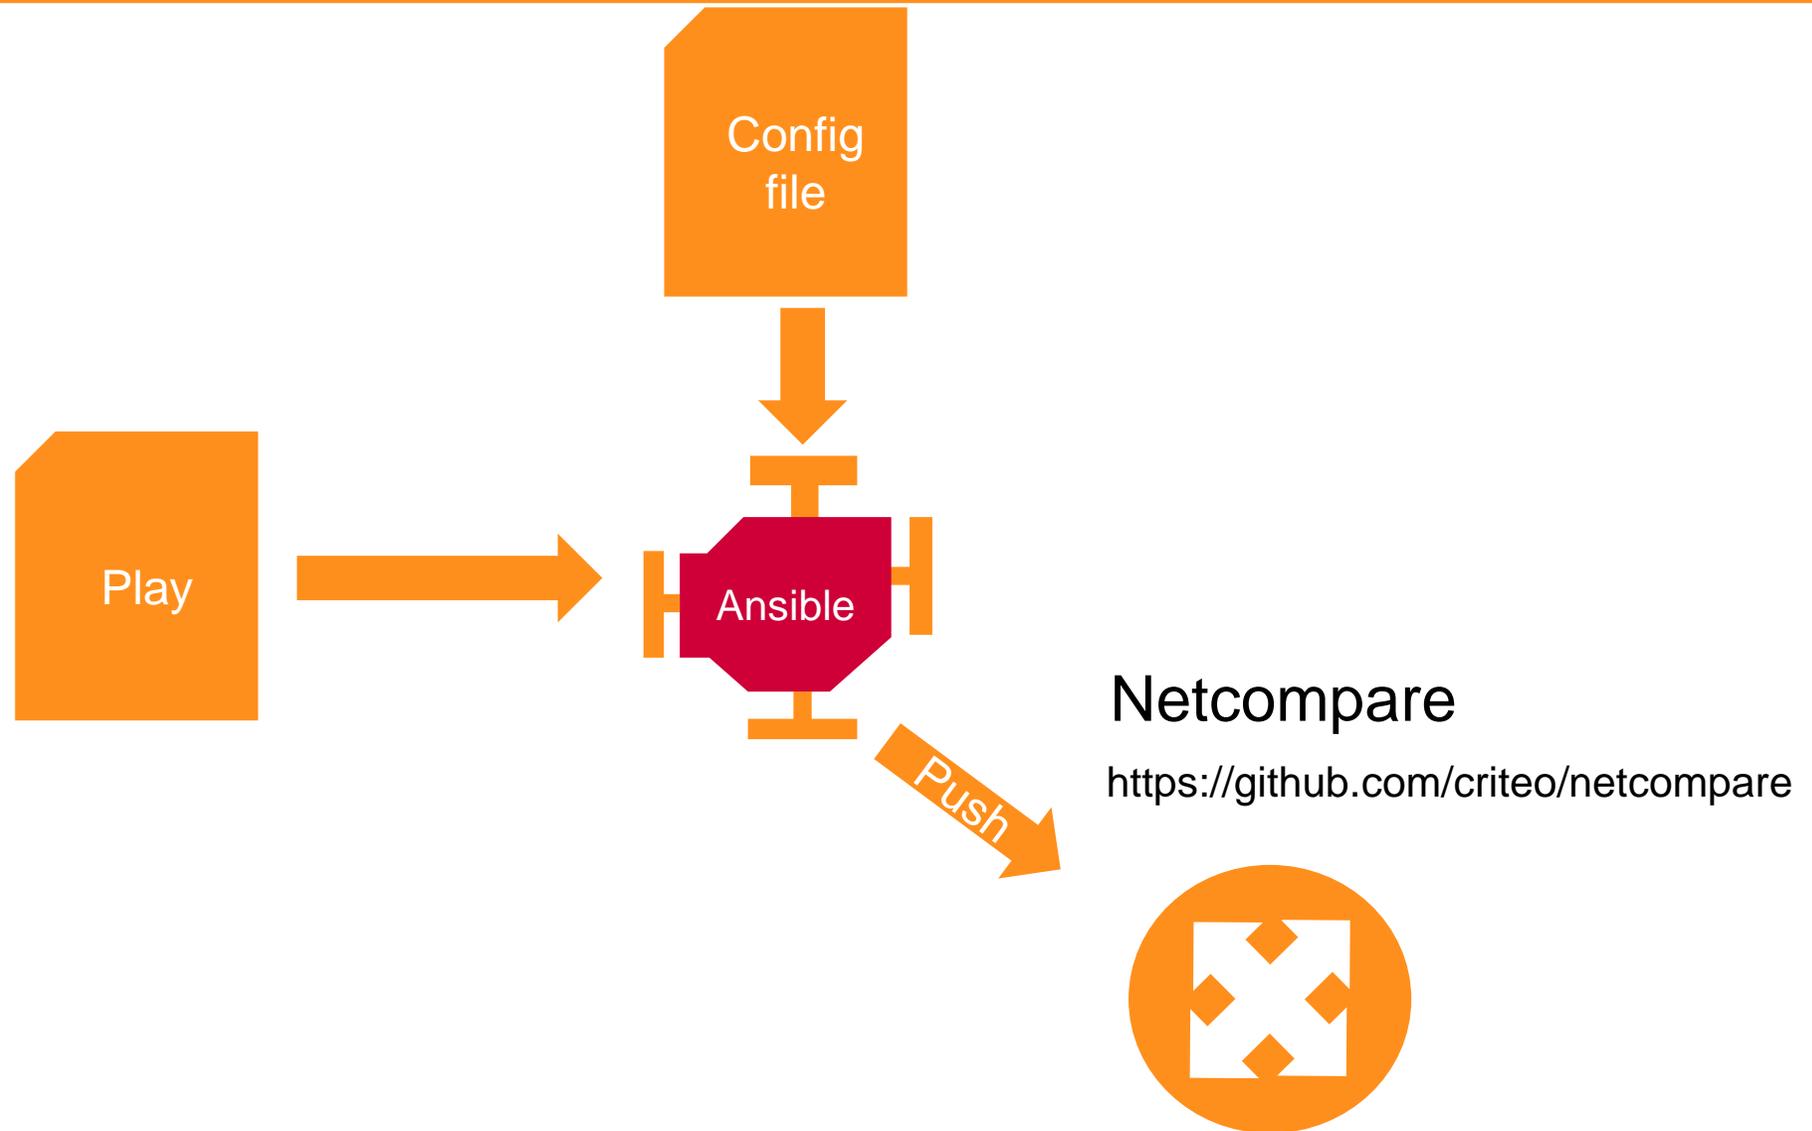
Automation is key



Automation is key



Automation is key



- Disaggregate, Distribute, Decentralize!
 - LBs, FWs, Switches
 - One OS – Multiple HW vendors
 - One template

- Onboard code:
- Counter streamers (kill SNMP)
 - Self healing devices
- Kill Centralization (LB/FW)
- On demand WAN TE

Join Us!

- Criteo is recruiting!
- Want to work with us on the future of DC networking?
- <http://www.criteo.com/careers/>
- <http://labs.criteo.com>
- rndrecruitment@criteo.com

- Come see us at the break!

Thank you.

Questions:

m.moriniaux@criteo.com

criteo.