



TURNING  
INTO REALITY

## Revisiting Routing Control Platforms with the Eyes and Muscles of Software-Defined Networking

Christian E. Rothenberg,  
Marcelo R. Nascimento,  
Marcos R. Salvador  
Telecomm. Research and  
Development Center (CPqD)  
Campinas - SP - Brazil  
esteve@cpqd.com.br

Carlos N. A. Corrêa,  
Sidney C. de Lucena  
Federal University of the Rio  
de Janeiro State (UniRio)  
Rio de Janeiro - RJ - Brazil  
carlos.correa@uniriotec.br

Robert Raszuk  
NTT MCL  
San Mateo - California - USA  
rr@nttmcl.com

ACM SIGCOMM HotSDN' 12 Workshop  
Helsinki, Finland, 13 August 2012

# Agenda

- Research in scope and contribution
- RouteFlow Control Platform
  - Multi-controller architecture
  - Proof of concept implementation: Single node abstract eBGP router
- Use Cases
- Challenges
- Conclusions and Future Work

## Research in scope and contribution

- Early work on Routing Control Platforms (RCP)  
[Ramjee 2006, Feamster 2004, Van der Merwe 2006, Wang 2009]
  - In operation at AT&T, considered a differentiator for "dynamic connectivity management".
- Research Question:
  - Re-examine the concept of RCP with the **visibility** (i.e., network-wide, multi-layer, flow and topology maps, full RIBs) and **direct control** capabilities (i.e., actual FIB installation, rich matching and instruction set) of the SDN abstraction set and the specifics of the OpenFlow choice
- RouteFlow glues virtualized IP routing stacks with OpenFlow
- RouteFlow acts as a new **indirection** layer for
  - routing protocol messages (e.g. BGP session terminates in servers)
  - RIB-(to-FIB)-to-OpenFlow transformations





# Software Defined IP Routing



Specialized Features

Specialized Control Plane

Specialized Hardware




Open interface



Open interface

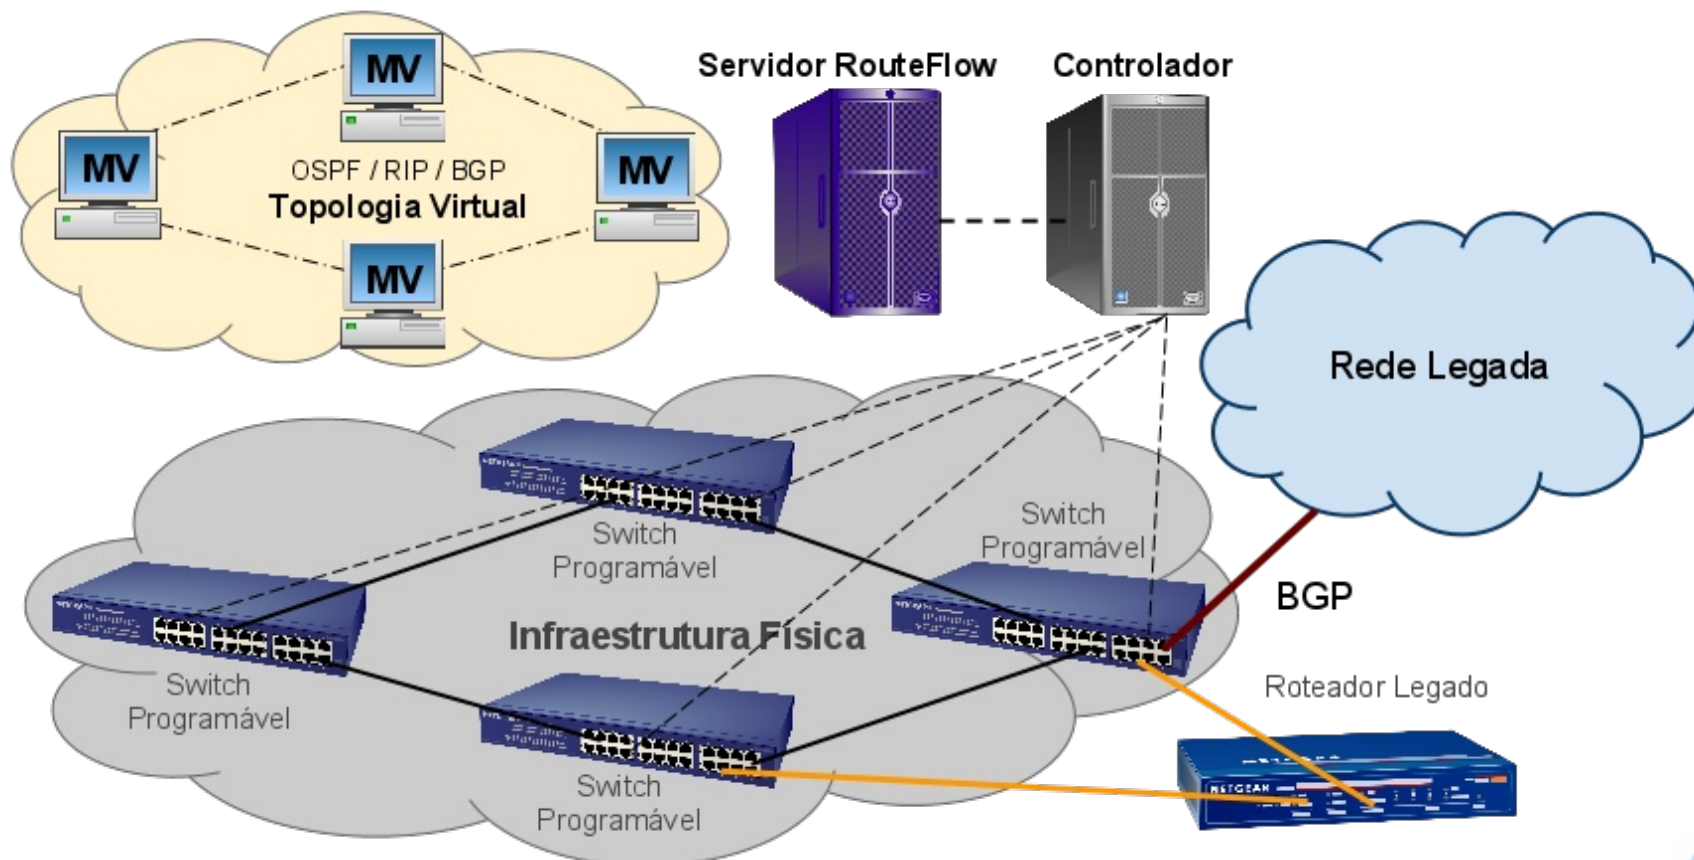


High cost  
Specialized config.  
Closed source  
Slow innovation



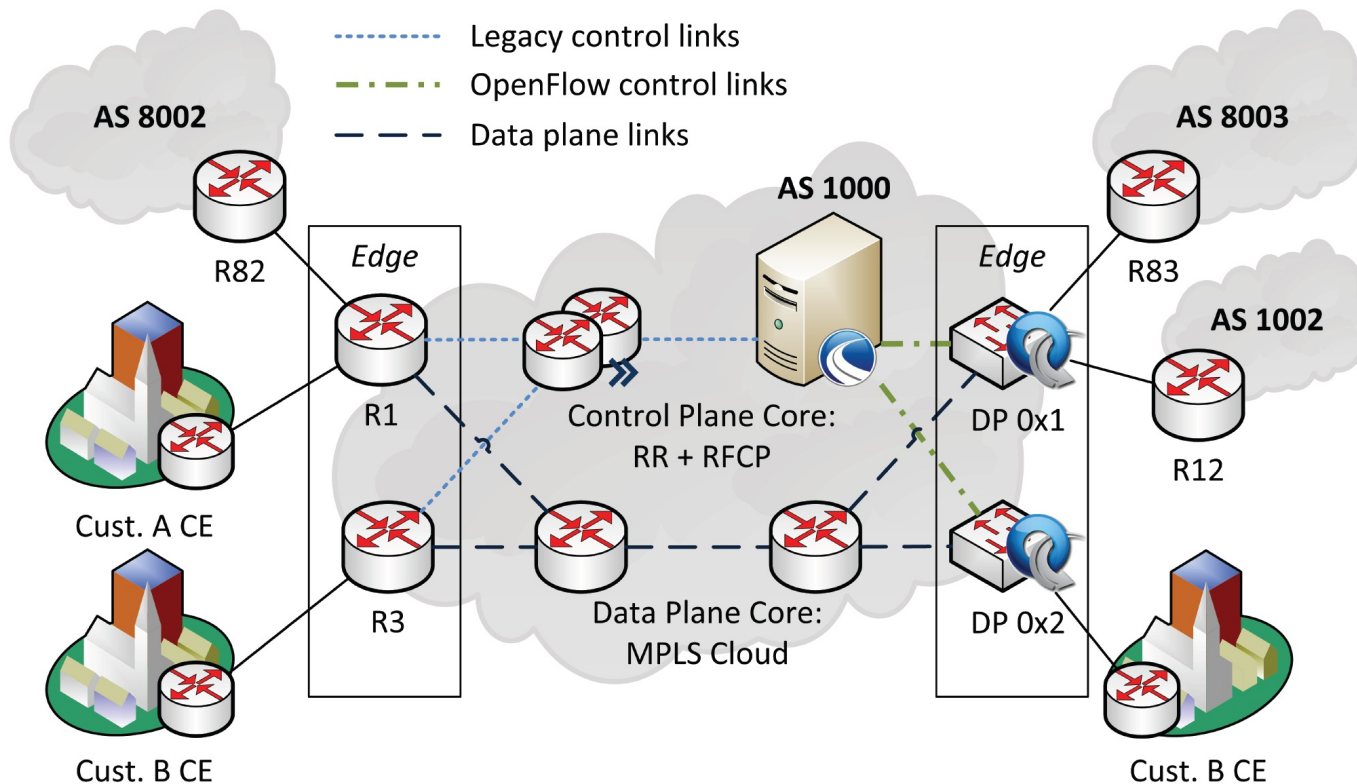
Source: McKeown

Low cost (commodity)  
Multi-vendor modularity  
Open source  
Rapid innovation




# Controller-Centric Hybrid Networking

- A migration path to roll out OpenFlow technology
- Not a revolution, but an evolution of current iBGP RRs to essentially eBGP Route Controllers
  - “BGP-free edge”: A cost-effective simplified edge for SW-driven innovations



# Design

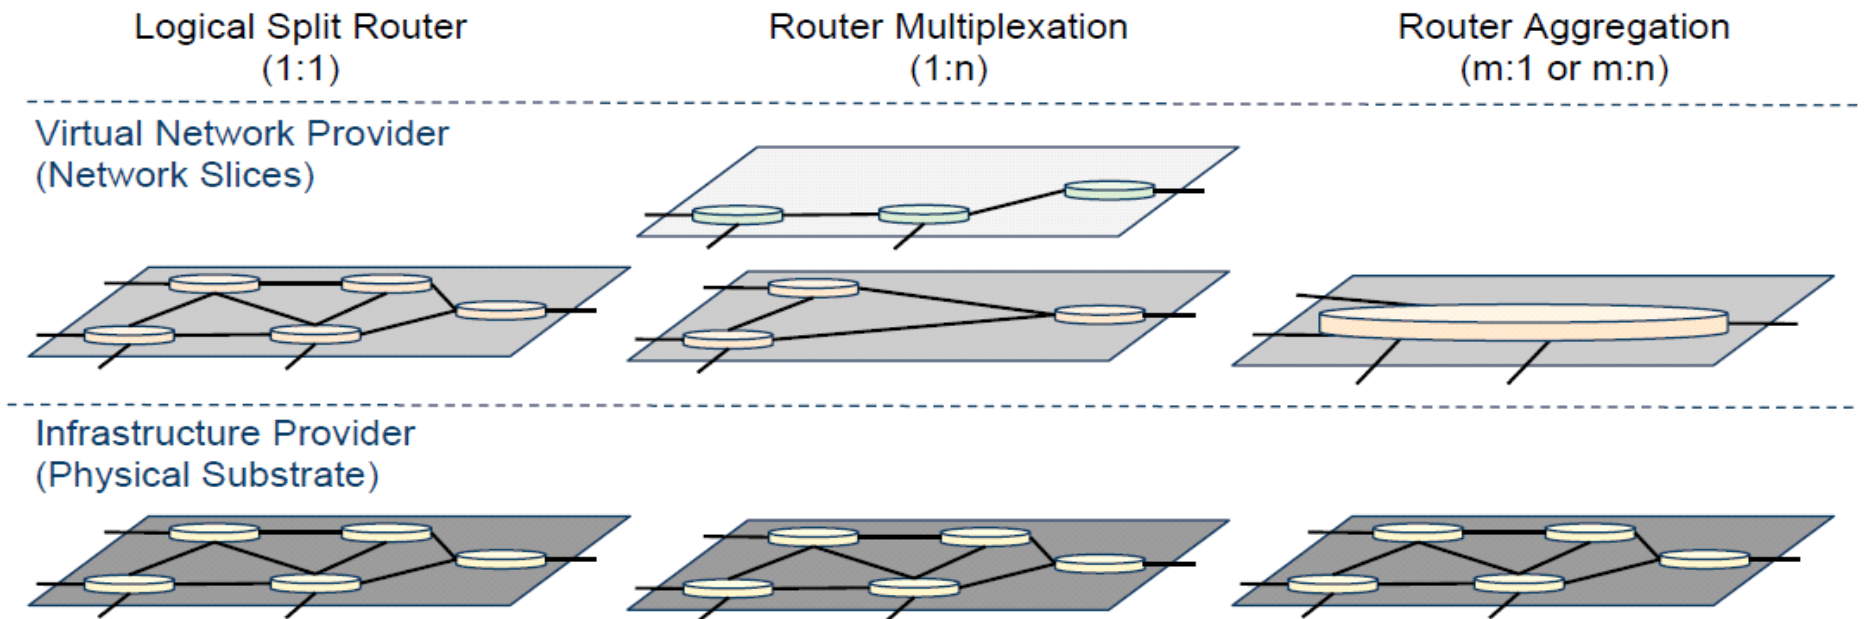


## Key Features

- Modular architecture
  - RF-Proxy
  - RF-Server
  - RF-Client
- Database layer
  - JSON-based IPC
  - Resilient core state
  - Programmer-friendly
- Multi-Controller support
  - NOX, POX, (Ryu)
  - Floodlight, Trema (planned)

# Modes of operation


- From logical routers (akin VRFs) to single node abstractions over flexible virtual networks.
- New design choices on the distribution of the control nodes.





# Aggregated BGP routing service

- Single node abstraction of a domain-wide eBGP router
  - Think modern multi-chassis routing architectures with external route processors and OpenFlow switches acting as line cards
- Aggregation logic defined in the RF-Server
- NOX, MongoDB, LXC



## Routing-centric use cases under research

- Engineered path selection
  - Think Google WAN, performance-based routing, etc.
- Optimal best path reflection
  - Per ingress/customer [draft-ietf-idr-bgp-optimal-route-reflection-01]
- Path protection with prefix independent convergence
  - Hierarchical FIBs w/ OF 1.X Tables + LFA route-precomputation
- Security
  - Data plane blackholes and middlebox injections,
  - Secure Inter-domain routing ideas (crypto intense S\*-BGP, etc..)
- Simplifying customer multi-homing
  - Easy to set and control cost/performance/policy-based routing
- IPv6 migration
  - Flow matching for service termination in v4-v6 migration solutions

# Challenges

- Centralized BGP
  - Shown to scale well in modern CPU architectures
  - Centralized does not mean not distributed (but removal from edge)
- Small OpenFlow table sizes
  - Transient limitation?
  - Expose existing FIB data structures as an IP lookup OF table?
  - Smart RIB&FIB reduction (e.g., simple [draft-ietf-grow-simple-va-04])
  - HW/SW flow offloading
- Limited OpenFlow processing in datapath
  - Transient / Un-optimized implementations
- High availability
  - Previous ideas from distributed RCPs
  - Database-centric designs
  - Development in-progress of “BGP SHIM” for transparent eBGP redundancy

## Conclusions

- RouteFlow is
  - a simple yet powerful (adaptable, inexpensive) routing architecture
  - a platform for real IP routing protocol experimentation
  - a tool for OpenFlow adoption via controller-centric hybrid networking
- Many open research questions and future work
  - OF 1.X, MPLS, OAM, GUI, policy languages, configuration mgm, etc.
- Opportunity for a community-driven development of competitive, deployable, open routing control solutions



**Evolving the IP routing landscape with OpenFlow/SDN**

Christian Esteve Rothenberg, PhD

Diretoria de Redes Convergentes (DRC)

esteve@cpqd.com.br

Thank you!

Questions?

[www.cpqd.com](http://www.cpqd.com)




## Live DEMO

- Access:
  - <http://go.cpqd.com.br/7API-demo>
  
- Indiana University GUI demo:
  - <http://goo.gl/T3Tqe>

# InCNTRE Deployment

- 4 Virtual routers
- 10 Gig and 1 Gig connections
- 2 BGP connections to external networks
- Remote Controller
- New User Interface



# Compare interfaces over the last 30 years

## “PC” user interfaces

```

Current date is Tue 1-01-1990
Enter new date:
Current time is 7:40:27.13
Enter new time:

The IBM Personal Computer BIOS
Version 1.10 (C)Copyright IBM Corp 1981, 1982

ROM BIOS
CIRQUAD CCM FERRAT CCM CHIBOK CCM SVS CCM DISKCOPY CCM
DISKCOPY CCM CIMP CCM DISKCOPY CCM MIRE CCM EPLM CCM
DISKOP CCM LINK CCM DISK CCM DISKOP CCM ART BAS
SAMPLES BAS MATHAGE BAS CALORAR BAS CALORAR BAS MISC BAS
EMORY BAS CIRCLE BAS MISCART BAS SPACE BAS DALL BAS
CMM BAS
20 Files(s)
AMIC command.com
COMMAND CCM 4350 5-07-82 12:00y
1 Files(s)
A>
  
```

## Network user interfaces

```

Router>en
Router#service-module gi/0 session
Trying 172.25.25.25, 2066 ... Open

Switch>en
Switch#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#int fas
Switch(config)#int fastEthernet 1/0/1
Switch(config-if)#ip address 10.10.10.2 255.255.255.0
Switch(config-if)#no shut
Switch(config-if)#end
Switch#
Switch#
Switch#
Building configuration...
[OK]
01:09:35: WSVS-S-CONFIG_I: Configured from console by console
Switch#
  
```



```

Router>en
Router#service-module gi/0 session
Trying 172.25.25.25, 2066 ... Open

Switch>en
Switch#conf t
Enter configuration commands, one per line. End with CNTL/Z.
Switch(config)#int fas
Switch(config)#int fastEthernet 1/0/1
Switch(config-if)#ip address 10.10.10.2 255.255.255.0
Switch(config-if)#no shut
Switch(config-if)#end
Switch#
Switch#
Switch#
Building configuration...
[OK]
01:09:35: WSVS-S-CONFIG_I: Configured from console by console
Switch#
  
```

Demystifying Configuration Challenges and Tradoffs in Network Based ISP Services (Benson, Akella, Shaikh SIGCOMM 2011)





# RouteFlow User Interface

- How to make network administration:
  - Simpler to implement
  - More robust and consistent
  - Easier to manage
- Automation and Abstraction
- Can you build very different interfaces with SDN backends?  
E.g., type: <http://netkarma.testlab.grnoc.iu.edu/rf/> or... <http://goo.gl/T3Tqe>

## RouteFlow Platform research topics

- High availability
- Integration of OF v1.1, v1.2 and v1.3
- LDP / MPLS support towards open-source LSR
- Realizing the northbound SDN abstractions
  - Specification / Configuration
  - Network Information Base
  - Knowledge Information Base
- Troubleshooting, testing, debugging, ...
- ...



## Collaborations and community developments

- Web-based UI & Internet 2 HW pilot [C. Small, Indiana] ✓
- Aggregated BGP Routing Service [C. Corrêa, Unirio] ✓
- SNMP plugin [J. Stringer, Google] ✓
- Optimal BGP best path reflection [R. Raszuk, NTT-MCL] ⌚
- OpenFlow v1.1 and v1.2 [w/ Ericsson] ⌚
- Open Label Switched Router [OSRF; Google] ⌚
- Multi-path, Fast-ReRoute, BGP-Sec, IPv6, ... [YOU?] ?



Open Source Routing




# ... building a community



Visits: 12,000+ (5,000+ Unique)

From over 1,100 cities of 90+ countries all over the globe!



<http://go.cpqd.com.br/routeflow/>

# 368

days since  
Project Launch



## Conclusions

- RouteFlow is
  - a simple yet powerful (adaptable, inexpensive) router design
  - a platform for real routing protocol experimentation
  - a tool for OpenFlow adoption via controller-centric hybrid networking
- Many open research questions
- Experimental research facilities are critical for validation
- Opportunity for a community-driven development of competitive, deployable, open routing control solutions



Open Source Routing

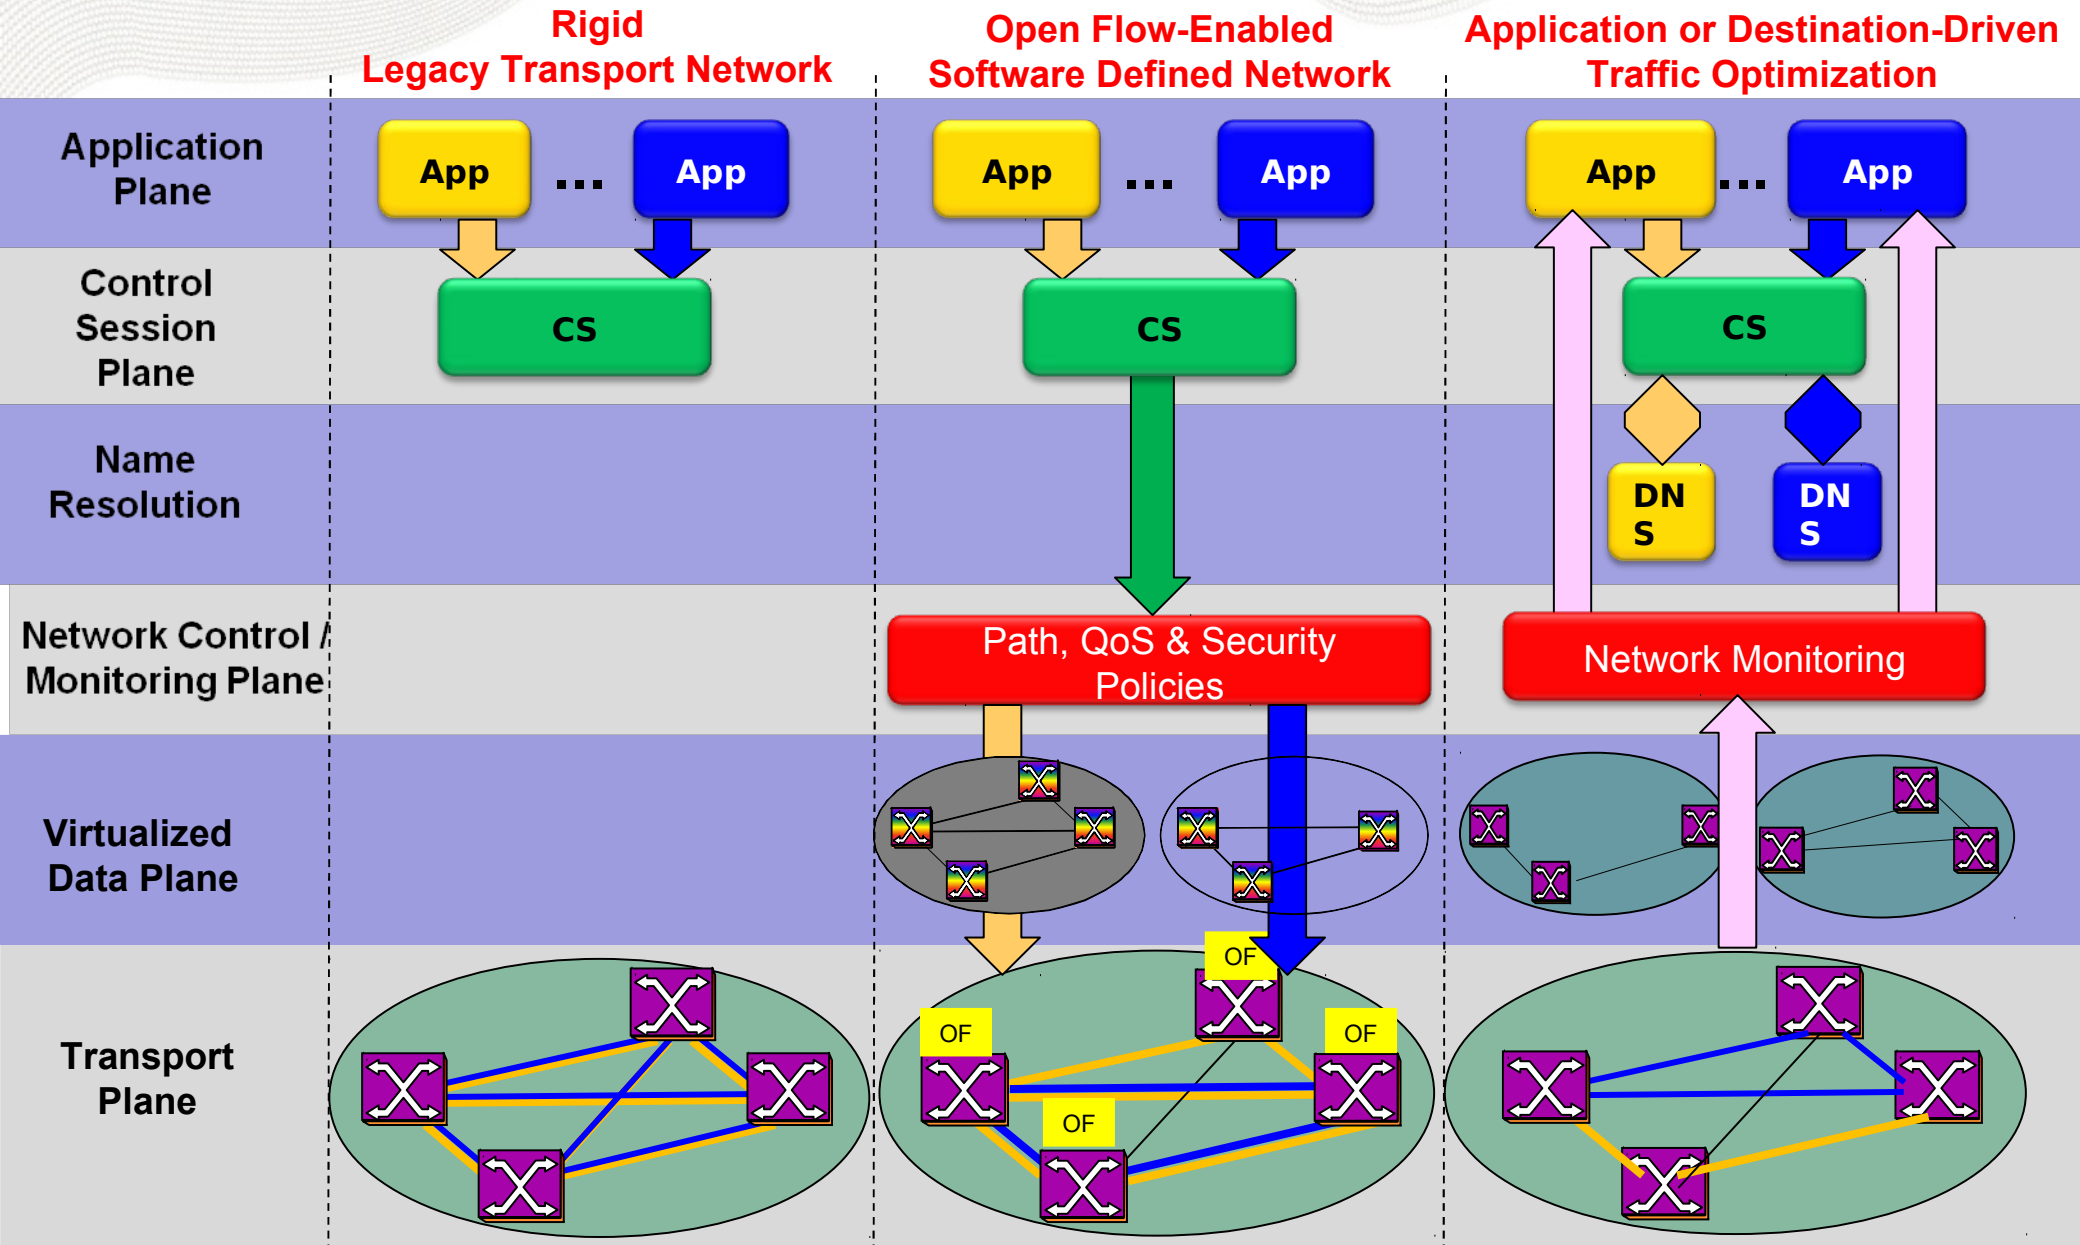


**Evolving the IP routing landscape with OpenFlow/SDN**

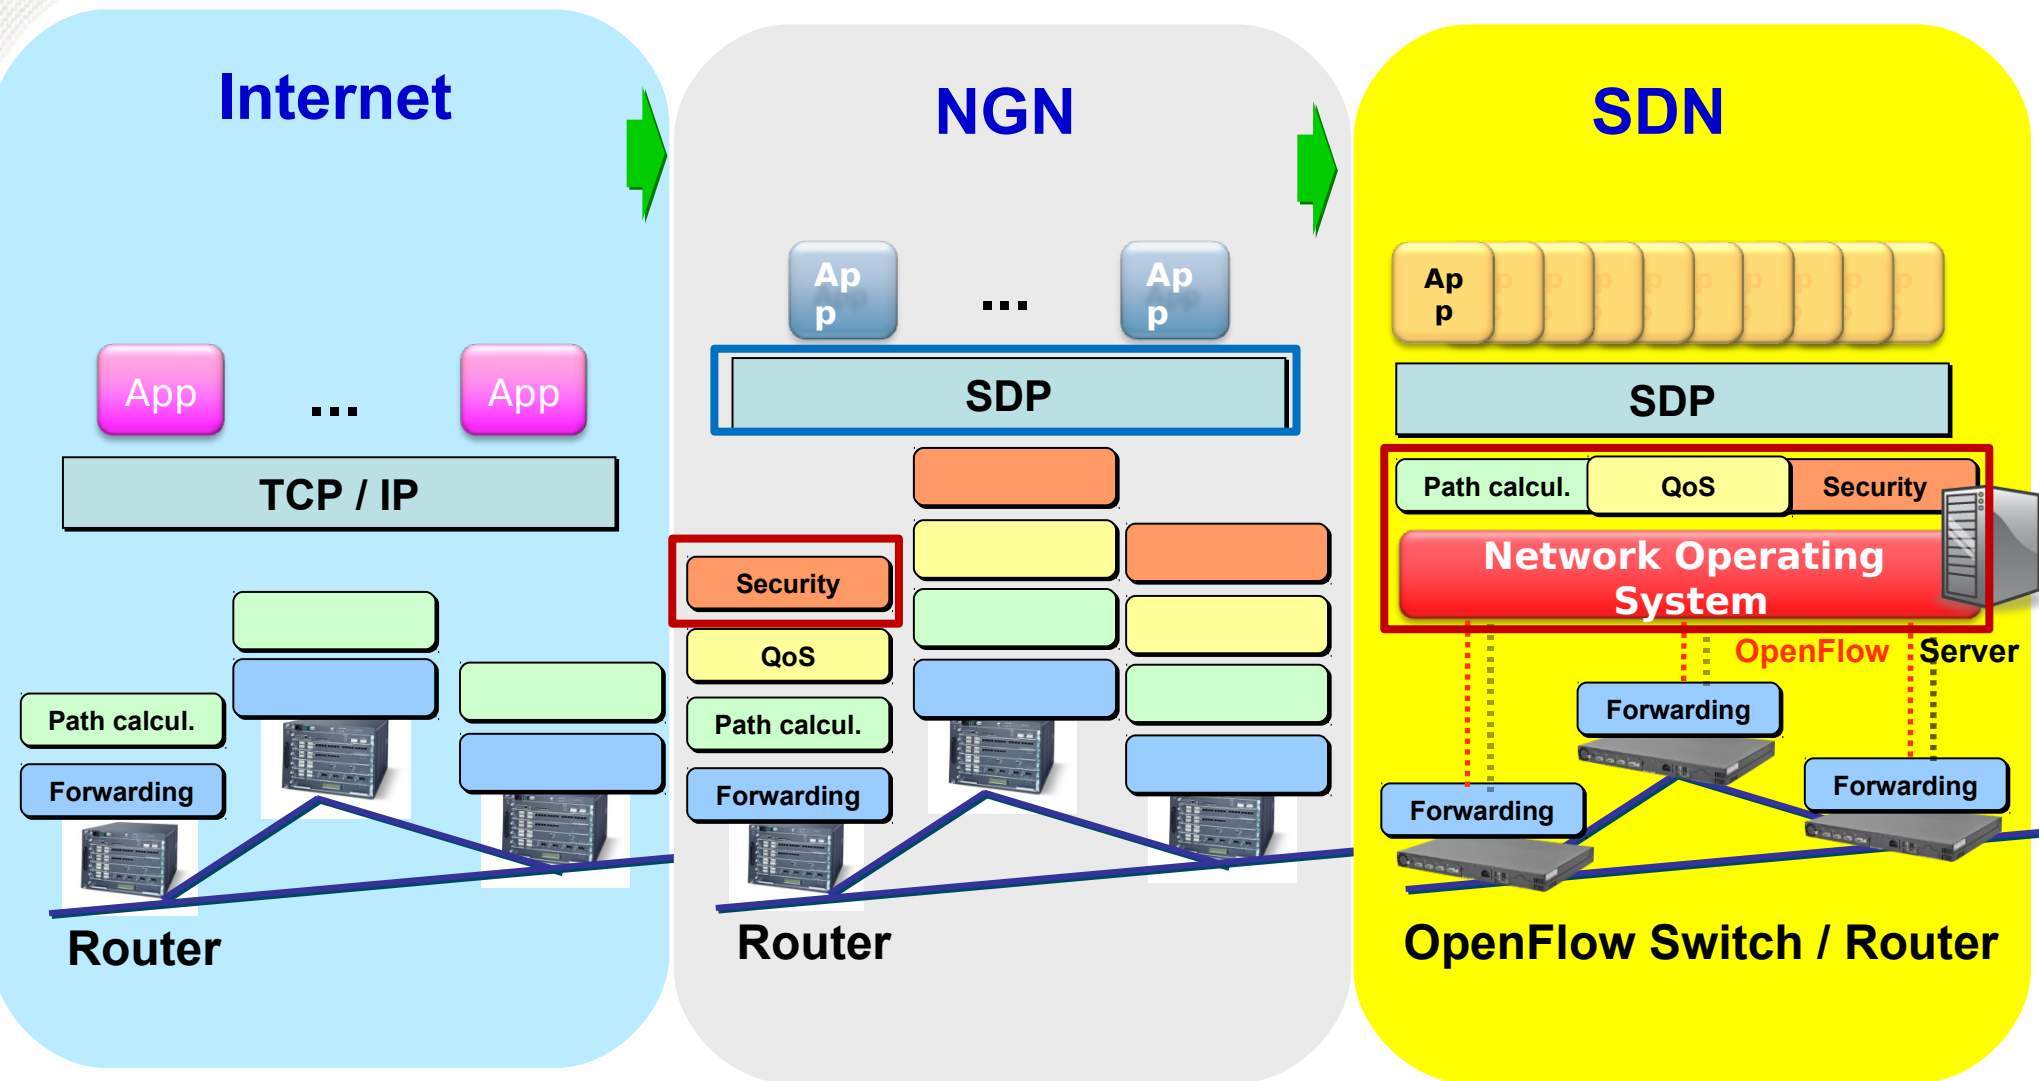
## Benefícios e impactos

- **Inovação tecnológica** em soluções de redes e serviços para os proprietários de infra-estrutura, os provedores de serviços e a comunidade de pesquisa.
- **Oportunidade para que empresas nacionais** possam competir e inovar na área de aplicações para gerenciamento e controle de redes de pacotes.
- **Novos modelos de negócio** que promovem redução de CAPEX e OPEX por meio de novos serviços (ex. alocação dinâmica de fatias/recursos da rede), reaproveitamento de ativos e automatização dos processos operacionais.
- **Diminuição do tempo ao mercado** na implementação de funcionalidades e soluções de redes integradas e customizadas à demanda do cliente.
- **Simplificação e barateamento** dos equipamentos pela diminuição dos requisitos mínimos de SW embarcado e pilhas de protocolos proprietárias.
- **Consolidação dos planos de controle e gerência** de infra-estruturas de rede, facilitando a convergência ampla e a migração para novos padrões e tecnologias de rede de transporte.

# SDN Converged Network Services

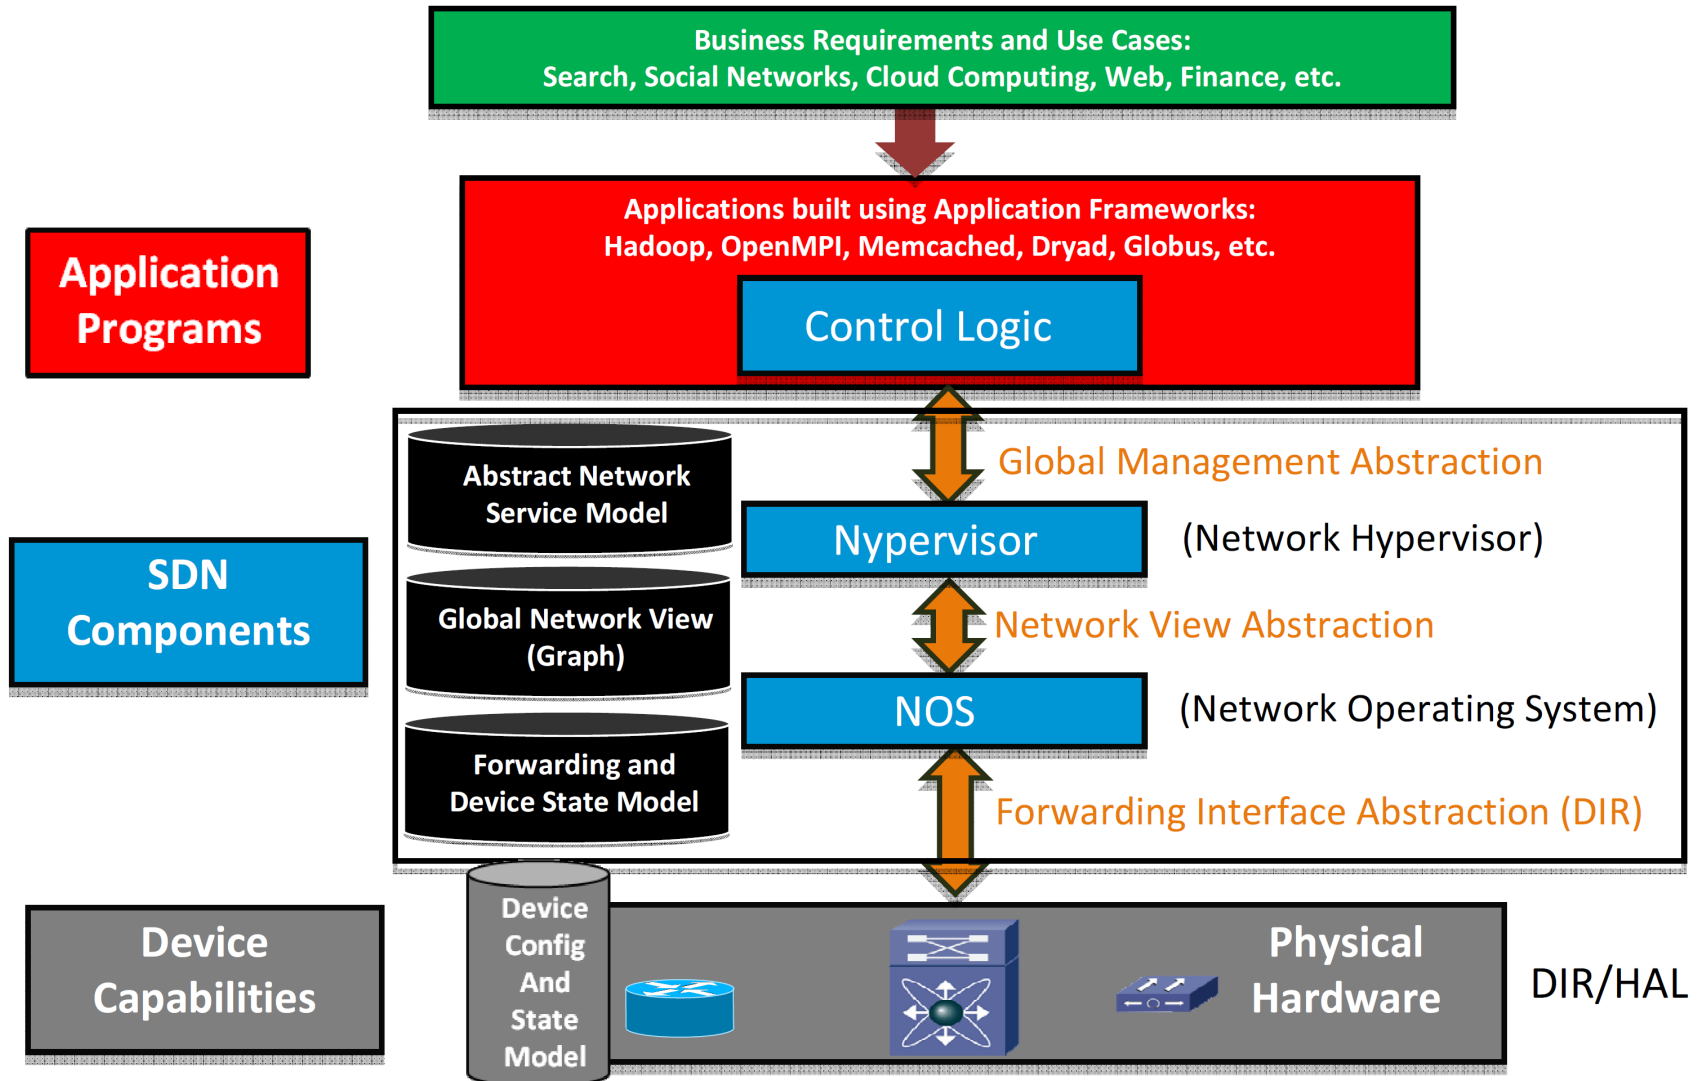


# Rede Convergente Definida por Software (RCDS)






# “Modern” OF/SDN architecture



# Control Plane Distribution Options



Data Path jointly controlled by standard on-box control plane and centralized off-box controller

Slide courtesy Frank Brockners

Legend:

Data plane	Control plane function