

OpenADN: Mobile Apps on Global Clouds Using Software Defined Networking



RAJ JAIN

Project Leader: Subharthi Paul
Washington University in Saint Louis
Saint Louis, MO 63130
Jain@cse.wustl.edu

Keynote at Mobile Cloud Computing & Services 2012
June 25, 2012, Low Wood Bay, Lake District, UK

These slides and audio/video recordings are available at:

http://www.cse.wustl.edu/~jain/talks/adn_mcs.htm

http://www.cse.wustl.edu/~jain/talks/adn_mcs.htm




1. Networking Application Trends
2. OpenFlow and SDN
3. OpenADN Vision and Extensions
4. Experimental Results
5. Key Features

2012: Where are we now?


- At the knee of Mobile Internet age (paradigm shift)
 - Computing (IBM 360) ⇒ Mini-computing (PDP11)
⇒ Personal Computing (Desktop, PC+MAC) ⇒ Laptops
⇒ Netbooks ⇒ Smart Phones + Tablets
Shift started on June 29, 2007 when iPhone was released.
- Most valued companies in the stock market are generally those that lead the paradigm shift
 - Automotive (General Motors) ⇒ Electrical (GE, Edison Electric) ⇒ Networking (Cisco + 3Com in 80's) ⇒ Internet (Netscape + Yahoo in 90's) ⇒ Mobile Internet (Apple +MS+ Google, 2010's)
- Note: Apple ≠ PC (MAC) company (mobile device company)
 - Google ≠ search engine (mobile device company)
- Also Social Networking (Facebook), Internet Retail (Amazon)

Internet Generations


- ❑ **Internet 1.0** (1969 – 1989)
 - ❑ Single ownership \Rightarrow Trust
 - ❑ complete knowledge
 - ❑ Algorithmic optimality \Rightarrow RIP



- ❑ **Internet 2.0** (1989–2009) Commerce
 - ❑ Multiple ownership of infrastructure \Rightarrow Distrust, **Security**
 - ❑ No knowledge of internal topology and resources
 - ❑ *Policy based* routing \Rightarrow BGP



- ❑ **Internet 3.0** (2009–2029) Commerce
 - ❑ Users, Content, Host ownership
 - ❑ Security, Mobility, Energy
 - ❑ Clouds, Services



Trend: Explosion of Mobile Apps and Clouds



- ❑ All top 50 Internet sites are services [Alexa]
- ❑ Almost all services are now mobile apps: Google, Facebook, Bank of America, ...
- ❑ Almost all services need to be global (World is flat)
- ❑ Almost all services use cloud computing (Easy management)

Networks need to support efficient service setup and delivery

Ref: Top 500 sites on the web, <http://www.alexa.com/topsites>


Washington University in St. Louis

http://www.cse.wustl.edu/~jain/talks/adn_mcs.htm


©2012 Raj Jain

Service Center Evolution

1. Single Server



2. Data Center



Load Balancers

SSL Off loaders

3. Global Clouds

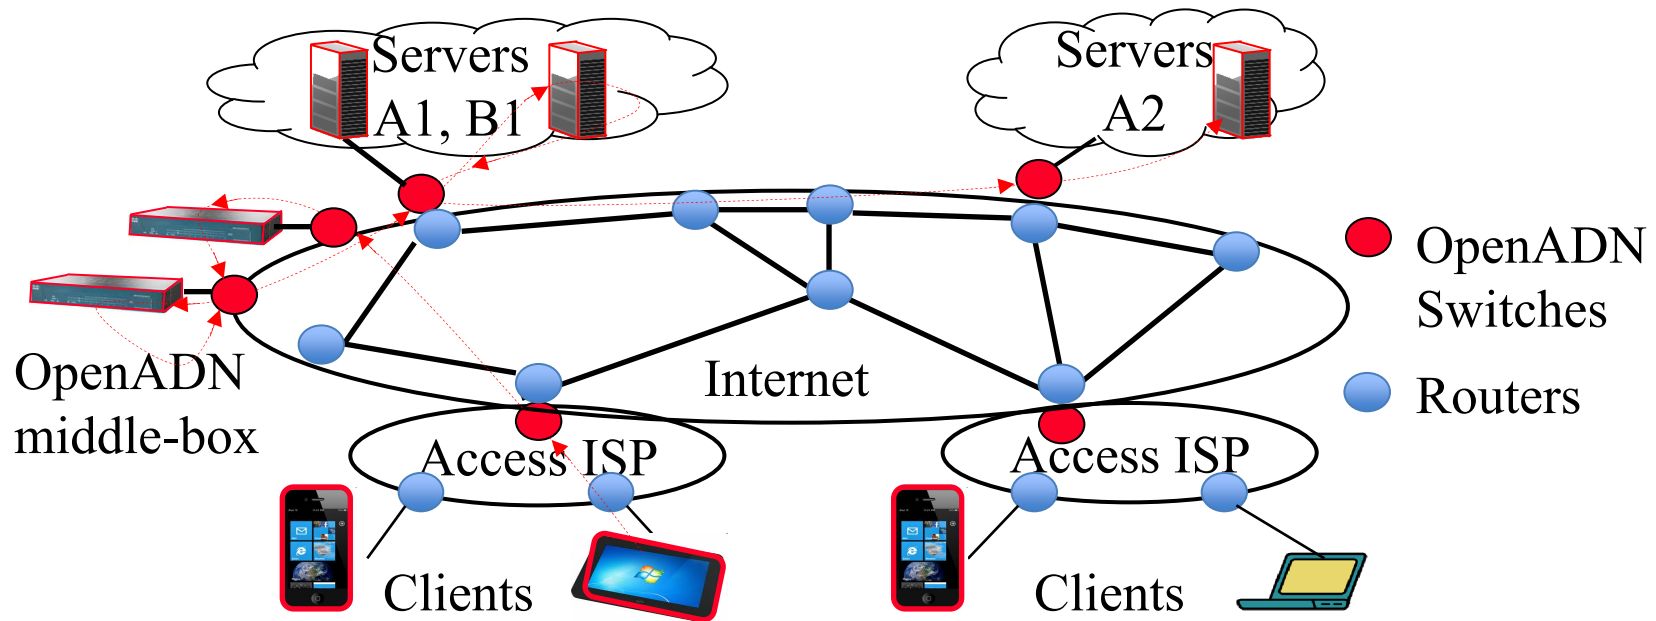


Global Internet


Need to make the global Internet look like a data center

Our Solution: OpenADN

- Open Application Delivery Networking Platform
Platform = OpenADN aware clients, servers, switches, and middle-boxes
- Allows Application Service Providers (ASPs) to quickly setup services on Internet using cloud computing



Google WAN




- ❑ OpenADN appliances are like Google appliances in Tier 3 ISPs
- ❑ Details of Google WAN are not public
- ❑ ISPs can not use it: L7 proxies require app msg reassembly

OpenADN Innovations

1. Cross-Layer Communication
2. MPLS like Labels
3. Extended OpenFlow flow-based handling, centralized policy control
4. Software Defined Networking: Standardized abstractions, Multi-Tenants, Control Plane programming for data plane
5. ID/Locator Split
6. Layer 7 Proxies without layer 7 visibility

Step 1: Separation of Control and Data Planes

- ❑ Control = Prepare forwarding table
- ❑ Data Plane: Forward using the table
- ❑ Forwarding table is prepared by a central controller
- ❑ Protocol between the controller and the forwarding element: **OpenFlow**
- ❑ Centralized control of policies
- ❑ Switches are simple.
Controller can be complex
Can use powerful CPUs
- ❑ Lots of cheap switches
= Good for large datacenters




Ref: [MCK08] "OpenFlow: Enabling Innovation in Campus Networks," OpenFlow Whitepaper, March 2008

<http://www.openflow.org/documents/openflow-wp-latest.pdf>

Washington University in St. Louis http://www.cse.wustl.edu/~jain/talks/adn_mcs.htm

©2012 Raj Jain


Centralized vs. Distributed



- Fully centralized is not scalable.
Fully distributed is not manageable.
⇒ Hierarchy

Step 2: Multi-Tenants Clouds


- ❑ Problem: Multiple tenants in the datacenter
- ❑ Solution: Use multiple controllers.
Each tenant can enforce its policies



- ❑ Significant industry interest \Rightarrow Open Networking Foundation, <https://www.opennetworking.org/>

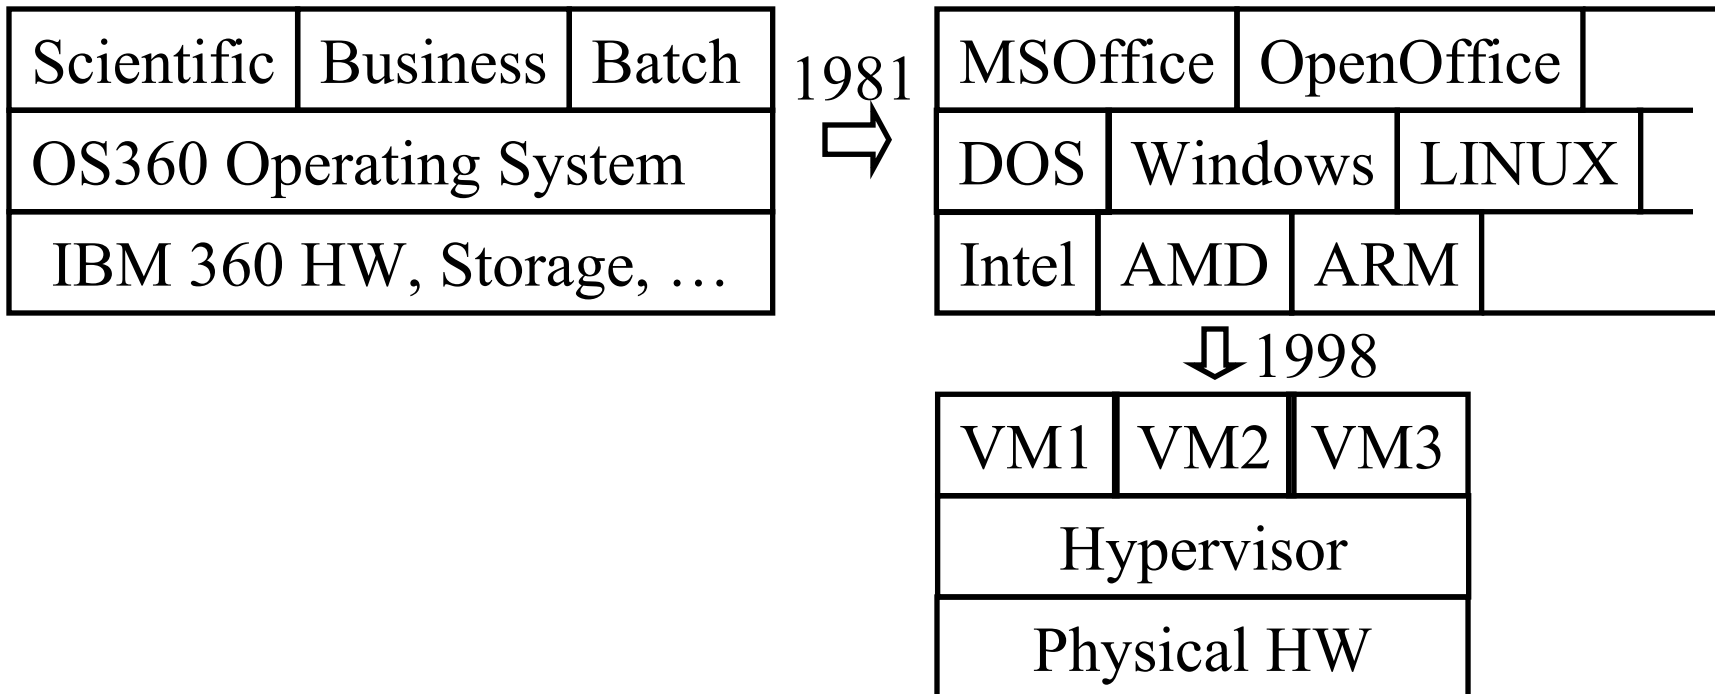
Step 3: Standardized Abstractions

- ❑ The routers are expensive because there is no standard implementation.
- ❑ Every vendor has its own hardware, operating/ management system, and proprietary protocol implementations.
- ❑ Similar to Mainframe era computers.
No cross platform operating systems (e.g., Windows) or cross platform applications (java programs).




Example: PC Paradigm Shift

- Computing became cheaper because of clear division of hardware, operating system, and application boundaries with well defined APIs between them
- Virtualization \Rightarrow simple management + multi-tenant isolation



Software Defined Networking

- Layered abstractions with standardized APIs




Ref: http://www.itc23.com/.../K1_McKeown-ITC_Keynote_Sept_2011.pdf

Washington University in St. Louis

http://www.cse.wustl.edu/~jain/talks/adn_mcs.htm

©2012 Raj Jain

SDN Architecture Component Examples



Ref: <https://courses.soe.ucsc.edu/courses/cmpe259/Fall11/01/pages/lectures/srini-sdn.pdf>

http://www.cse.wustl.edu/~jain/talks/adn_mcs.htm


SDN Impact

- ❑ Why so much industry interest?
 - ❑ Commodity hardware
 - ⇒ Lots of cheap forwarding engines ⇒ Low cost
 - ❑ Programmability ⇒ Customization
 - ❑ Sharing with Isolation ⇒ Networking utility
 - ❑ Those who buy routers, e.g., Google, Amazon, Docomo, DT will benefit significantly
- ❑ Opens up ways for new innovations
 - ❑ Dynamic topology control: Turn switches on/off depending upon the load and traffic locality
 - ⇒ “Energy proportional networking”


Before



After



Life Cycles of Technologies



Industry Growth: Formula for Success



Innovators

⇒ Startups

⇒ Technology


Differentiation

Big Companies

Manufacturing


⇒ Price differentiation

Number of
Companies



- ❑ Paradigm Shifts ⇒ Leadership Shift
- ❑ Old market leaders stick to old paradigm and loose
- ❑ Mini Computers → PC, Phone → Smart Phone, PC → Smart Phone

OpenADN in SDN's Layered Abstraction



- SDN provides standardized mechanisms for distribution of control information


Application vs. Network Flow Classes

- Network-level Policies:
 - Specified by ISPs, e.g., routing, traffic engineering, congestion control, ...
 - Applied to all packets that belong to a network flow Class, e.g., by source-destination addresses and MPLS tags
 - Easily enforced by routers.
Packet header is sufficient to enforce network-level policies
- Application-level Policies:
 - Specified by ASPs, e.g., Send all voice + video messages to Server group 1, accounting messages to Server 2
 - Bill advertising data to ASP and not client
 - Enforced by middle boxes
 - Usual way: Look at application messages,
Requires message reassembly and decryption in middle boxes.
 - Our solution: Use cross-layer communication

OpenADN Features

Message level:

- Server selection
- Load balancing between servers
- Fault tolerance
- Server mobility
- User Mobility
- Secure L5-L7 headers and data
- Middlebox services: Intrusion detection, Content based routers, application firewalls, ...
 - Control plane and data plane MBs
- Middlebox traversal sequence
- Message level policies
- TCP Splicing



Networking: Failures vs Successes

- ❑ 1986: MAP/TOP (vs Ethernet)
- ❑ 1988: OSI (vs TCP/IP)
- ❑ 1991: DQDB
- ❑ 1994: CMIP (vs SNMP)
- ❑ 1995: FDDI (vs Ethernet)
- ❑ 1996: 100BASE-VG or AnyLan (vs Ethernet)
- ❑ 1997: ATM to Desktop (vs Ethernet)
- ❑ 1998: ATM Switches (vs IP routers)
- ❑ 1998: MPOA (vs MPLS)
- ❑ 1999: Token Rings (vs Ethernet)
- ❑ 2003: HomeRF (vs WiFi)
- ❑ 2007: Resilient Packet Ring (vs Carrier Ethernet)
- ❑ IntServ, DiffServ, ...



Technology alone does not mean success.

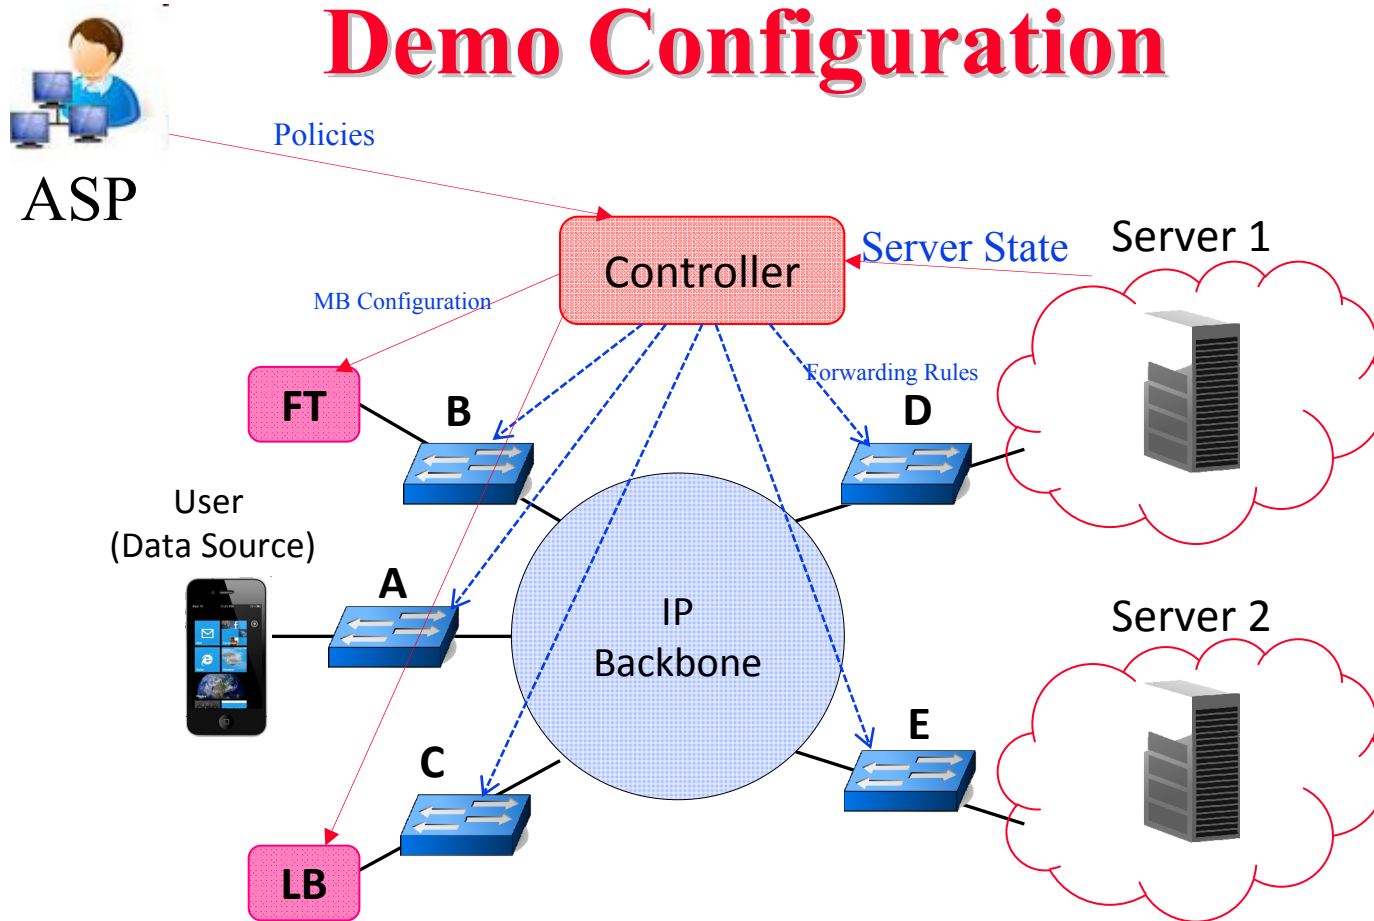
Key Features of OpenADN

1. Edge devices only.
Core network can be current TCP/IP based, OpenFlow or future SDN based
2. Coexistence (Backward compatibility):
Old on New. New on Old
3. Incremental Deployment
4. Economic Incentive for first adopters
5. Resource owners (ISPs) keep complete control over their resources





**Most versions of Ethernet followed these principles.
Many versions of IP did not.**

Demo Configuration



- ❑ Single user and single ASP with 2 servers
- ❑ OpenADN Appliances: A, B, C, D, E
- ❑ ISP offers ADN services: Fault tolerance and Load Balancing

Validation of Functionality




Resource Control

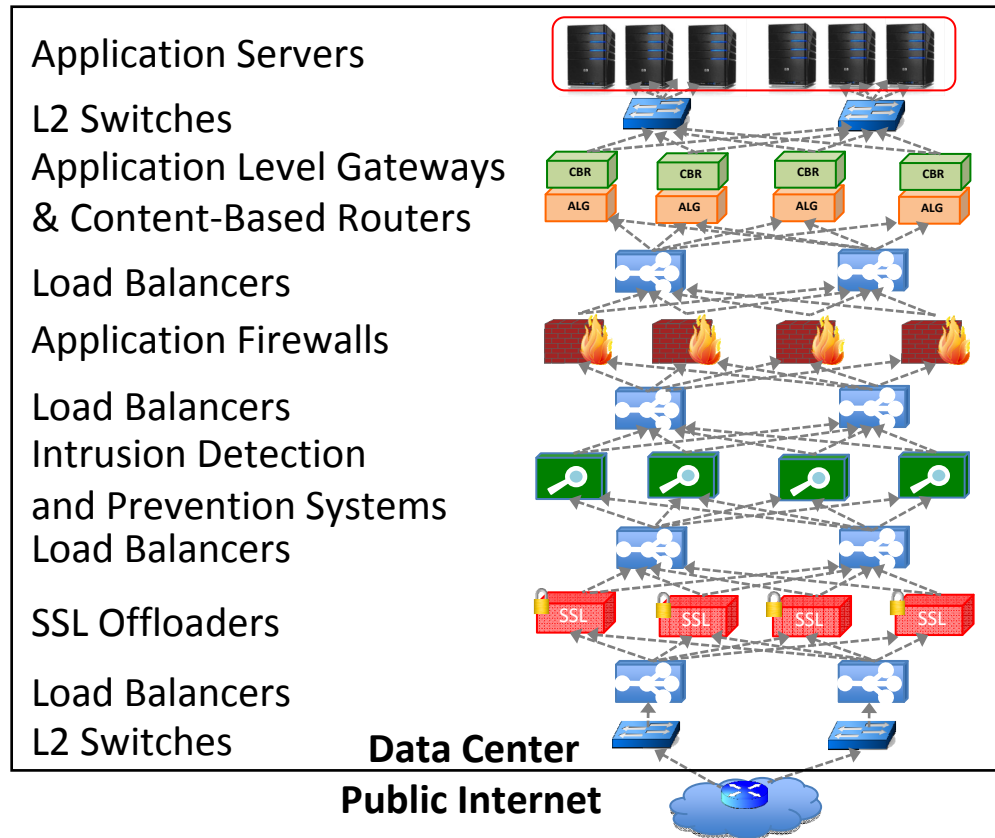
- ❑ ASPs keep complete control of their data.
ISP does not have to look at the application headers or data to enforce application level policies
- ❑ ISPs keep complete control of their equipment.
ASPs communicate their policies to ISP's control plane
- ❑ Middle boxes can be located anywhere on the global Internet
(Of course, performance is best when they are close by)
- ❑ ISPs own OpenADN switches and offer them as a service
- ❑ ASPs or ISPs can own OpenADN middle boxes
- ❑ No changes to the core Internet

Beneficiaries of This Technology

- ❑ Equipment/Software vendors: OpenADN-aware appliances
- ❑ ASPs: Deploy servers anywhere and move them anytime
- ❑ ISPs: Offer new application delivery/middlebox services
- ❑ Cloud Service Providers (CSPs): Freedom to move VMs, Less impact of downtime
- ❑ CDNs, e.g., Akamai, can extend into application delivery

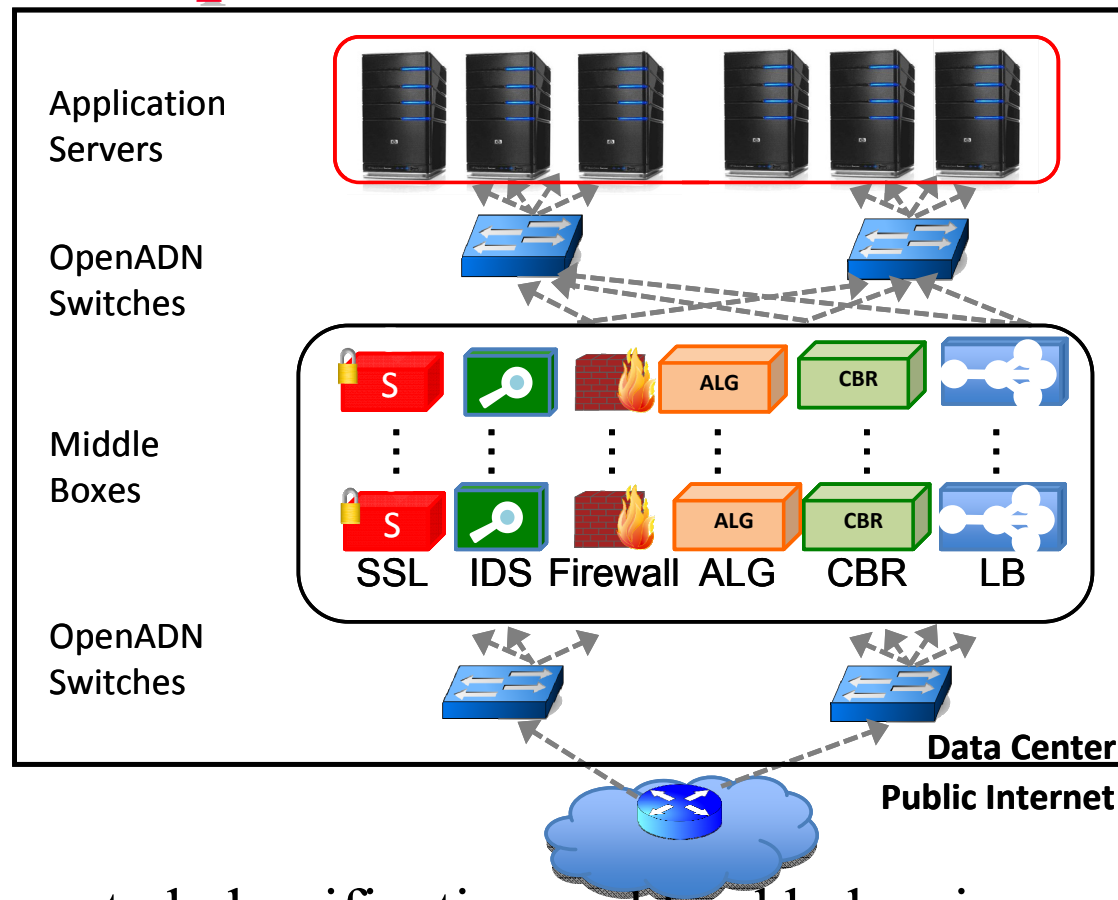


Data Center Applications








- ❑ Repeated classification and load balancing
- ❑ No application level control over MBs traversed
- ❑ Unnecessary traversals and reduced performance

OpenADN in Data Center



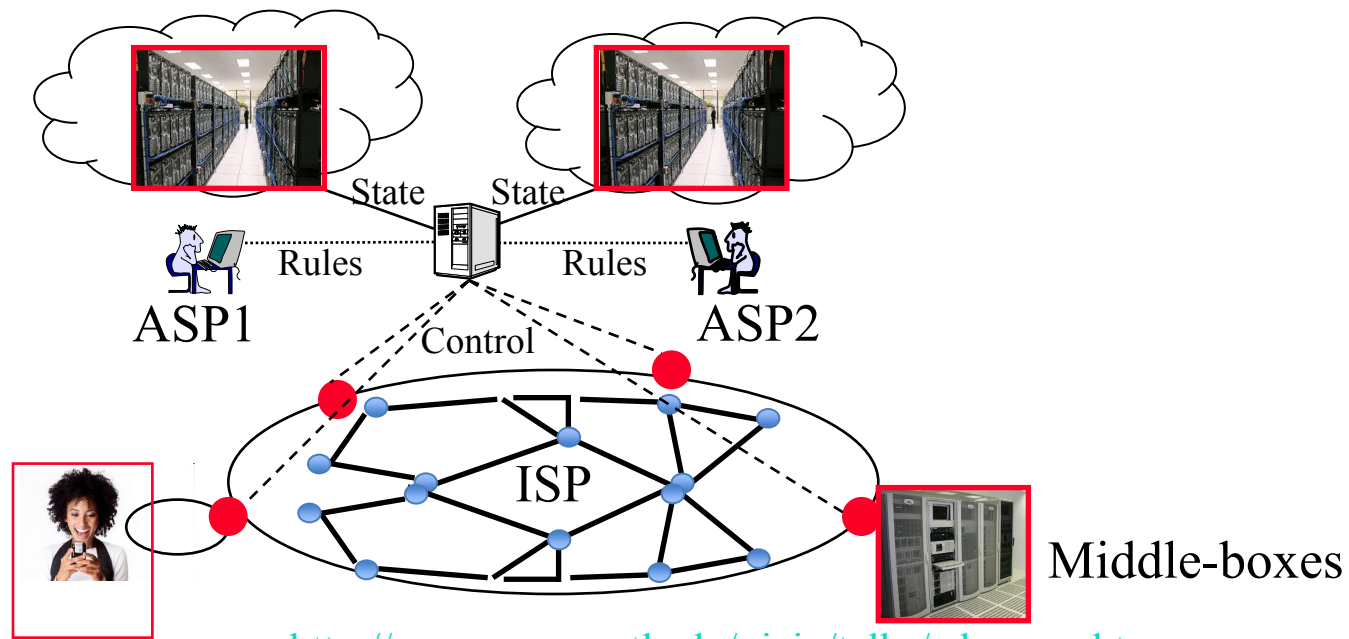
- ❑ No repeated classification and load balancing
- ❑ Application flow specific traversal through MBs
- ❑ Reduced number of appliances and increased performance


The Narrow Waist



OpenADN Without Standard OpenFlow/SDN

- ❑ OpenADN clients, servers, middle-boxes use only APLS labels.
- ❑ OpenADN aware devices need an API to communicate with controllers
- ❑ API can be vendor specific





Summary

1. Knee of **mobile internet** paradigm shift
Explosion of Apps using cloud services
2. OpenADN appliances can provide ASPs networking services they need
3. OpenADN extends using best of OpenFlow, SDN, MPLS, ID/Locator Split, Cross-layer communications, middle box appliances
4. Keeps resource control under resource owners.
Can be implemented incrementally now
5. Trend is towards simplifying and standardizing router interfaces \Rightarrow Software defined networking

Application Delivery: Opportunity for ISP's

Thank You!

A cartoon illustration of a man with dark, curly hair, a large nose, and a slight smile. He is wearing a light blue short-sleeved shirt, a pink tie, and dark trousers with a brown belt. He is holding a wooden signpost that supports a large white rectangular sign with a black border. The sign contains the word "Questions?" in a large, black, serif font. The signpost is also visible at the top of the sign.

Questions?