UNIVERSITY OF KENTUCKY®

A Simplified SDN-Driven All-Campus Science DMZ

Jacob Chappell, C. Lowell Pike, Dr. Cody Bumgardner, Dr. Brent Seales, Dr. James Griffioen

University of Kentucky

UK UNIVERSITY OF KENTUCKY

, 2018

KFC

Where is Kentucky?

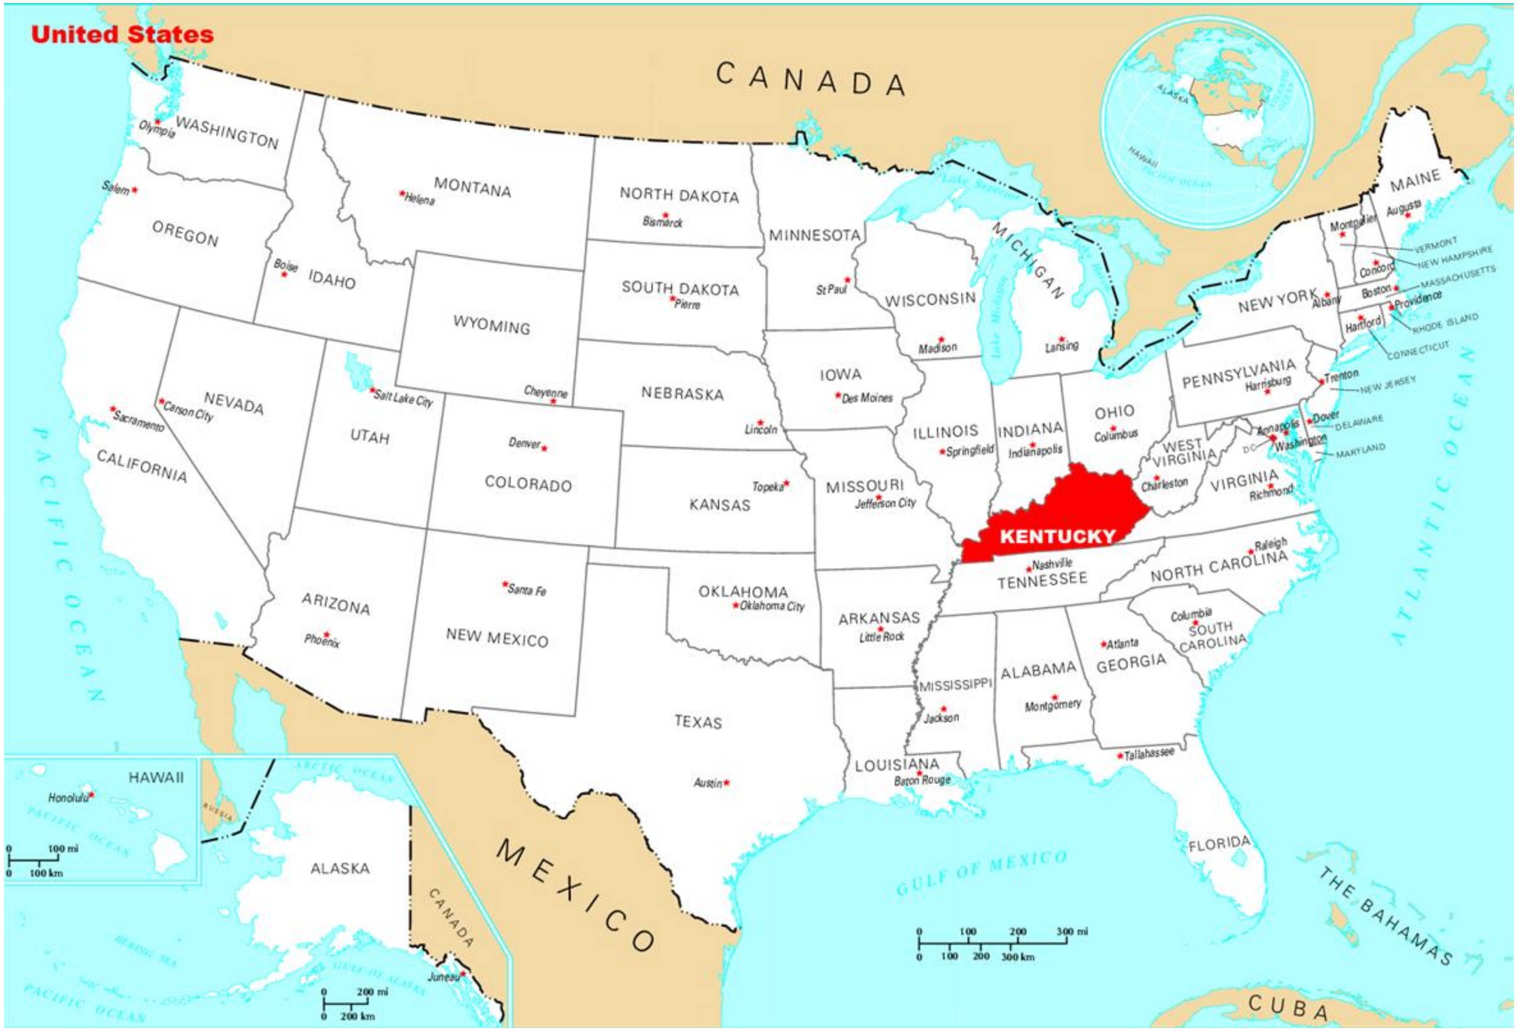

Horse Racing and Breeding

UNIVERSITY OF KENTUCKY

Bourbon Whiskey

UK
UNIVERSITY OF KENTUCKY

Tobacco

UNIVERSITY OF KENTUCKY

Agenda

- Big data woes on the campus network
- Standard science DMZ solution
- Brief SDN overview
- A new DMZ approach
- Some results

Big Data in Research

- Large data sets are becoming increasingly prevalent in research.
  - Machine Learning
  - Data Mining
  - Analytics
  - Modeling
  - Visualization
  - Simulation
  - …
- Furthermore, researchers often need to move their large datasets between research sites and into and out of cloud storage.
- Traditional campus networks are not designed to support pervasive big data usage.

UNIVERSITY OF KENTUCKY

Image Source: http://ohno.es/67076403.jpg

Typical Campus Network

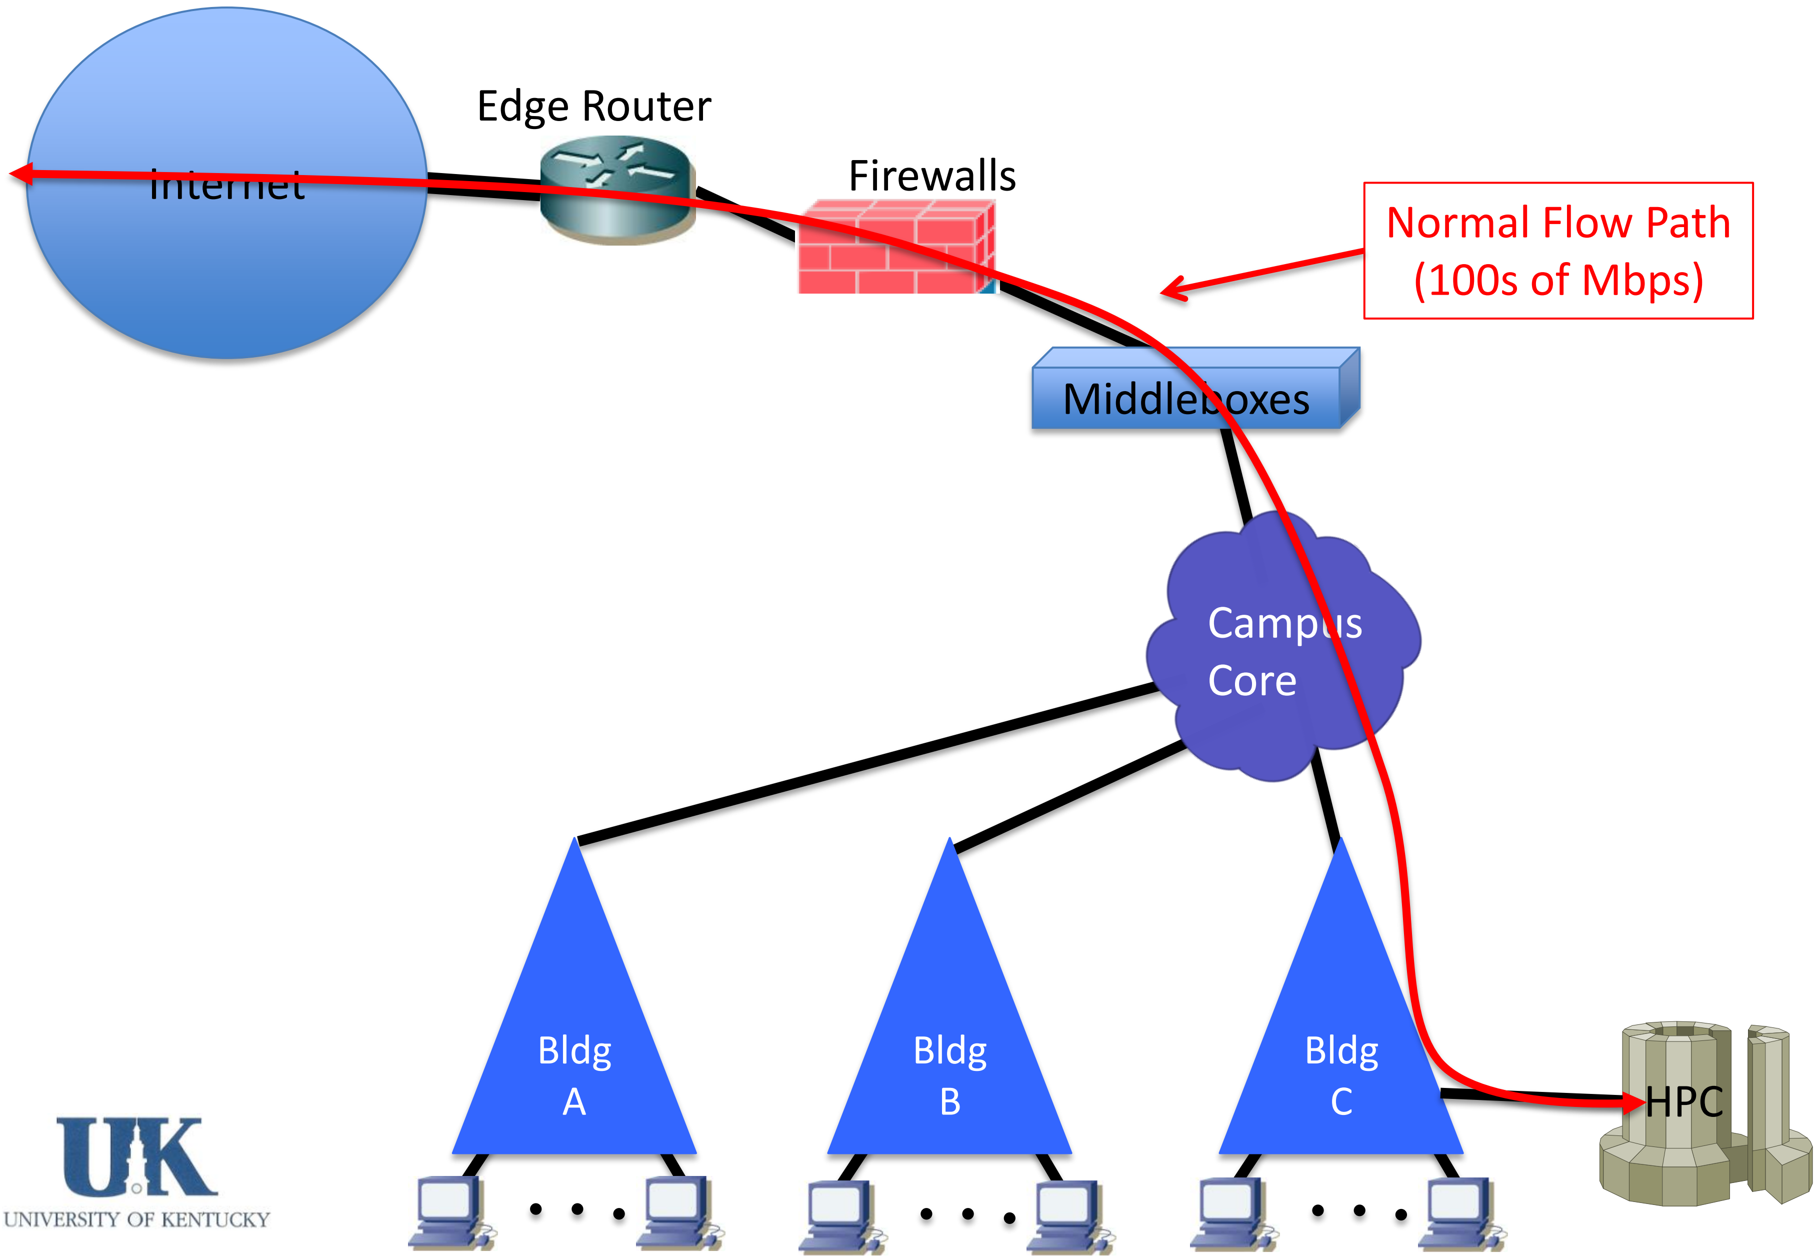

Normal Flow Path (100s of Mbps)

Typical Campus Network

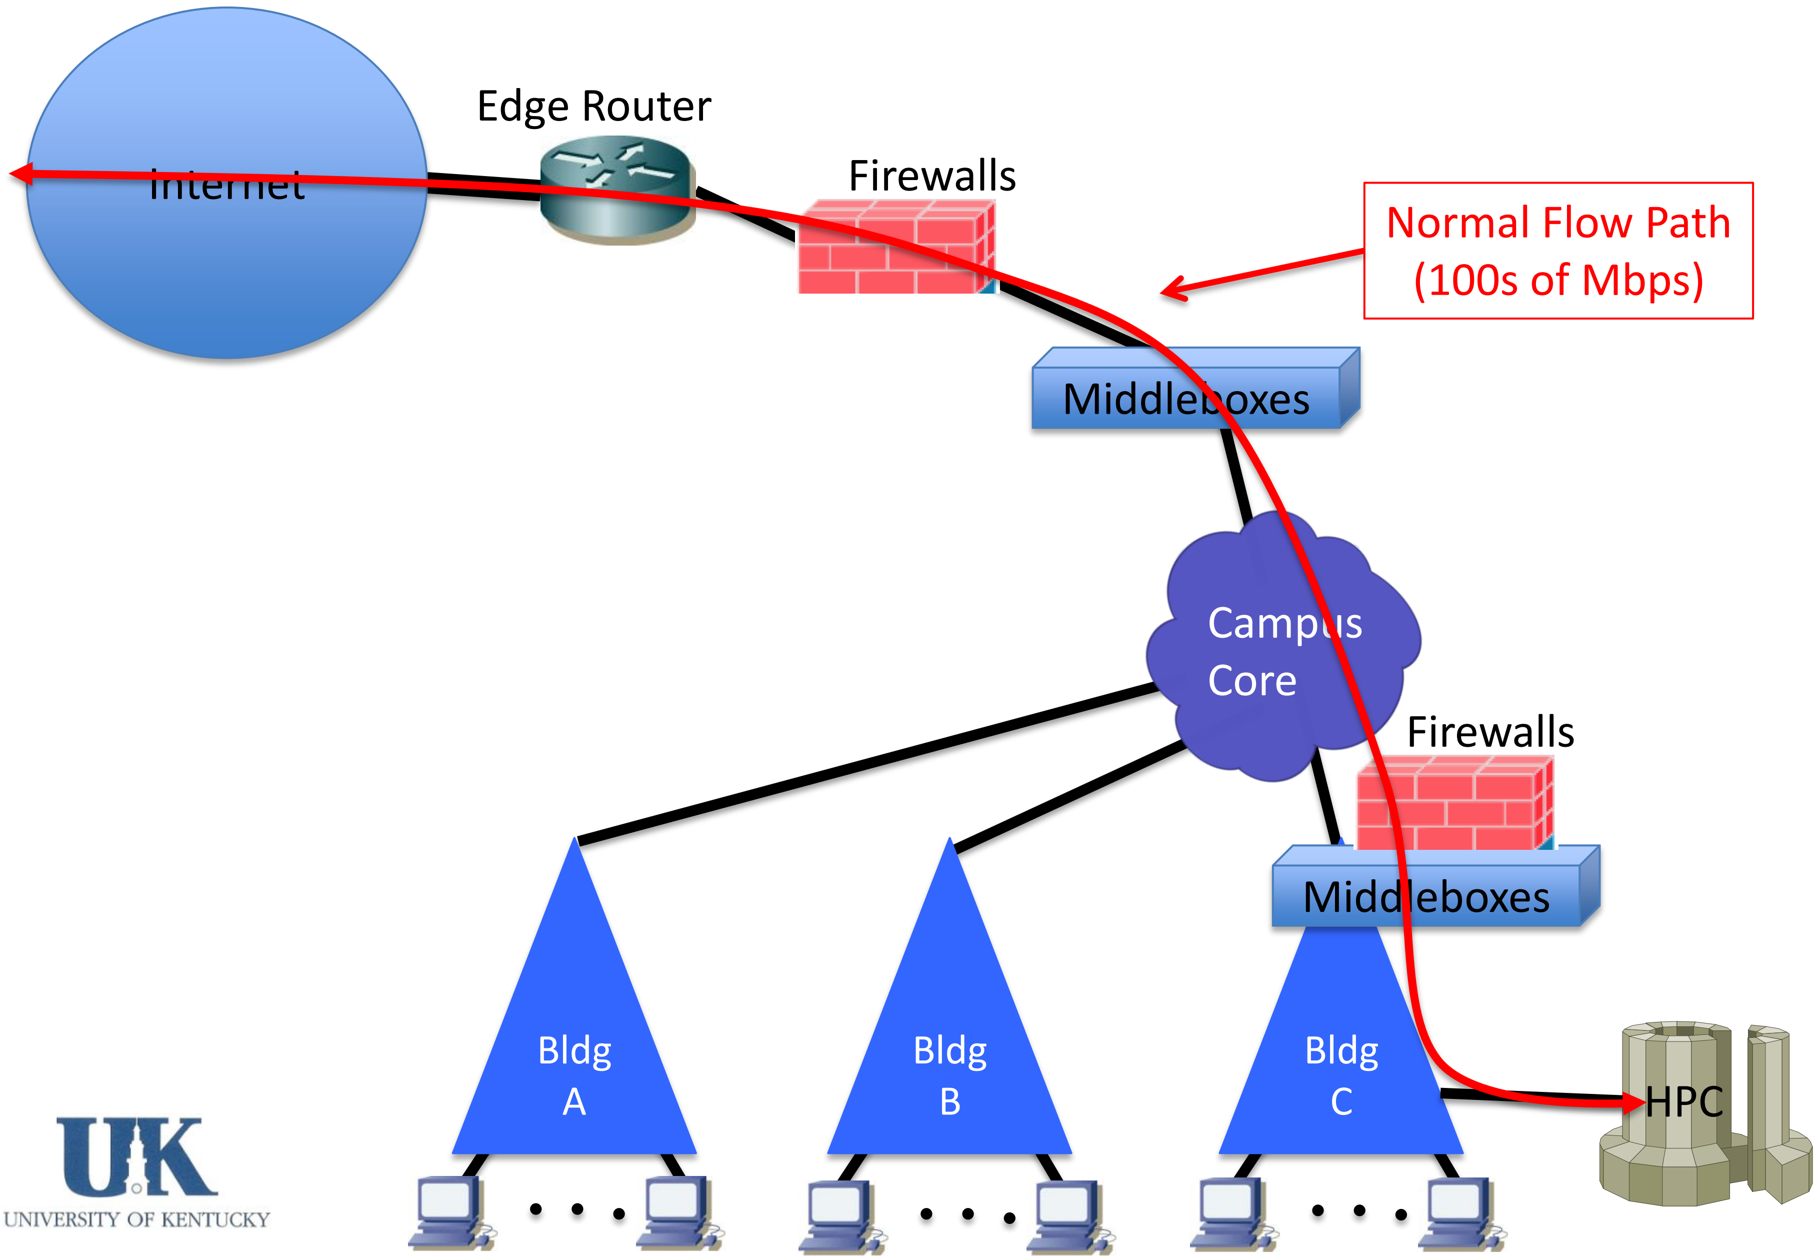

Big Data Woes on Campus Network

- Middleboxes
- Competition: 45K students, faculty, staff
- Refresh needed: older infrastructure
- Backpressure: even with upgrades

Middleboxes

- Packet inspecting/modifying devices scattered throughout the campus network.
- Provide important services essential to a stable and secure campus network.
- Impose intentional and unintentional bottlenecks in network performance.
- Provided services include:
  - Network Address Translation (NAT)
  - Intrusion Detection (e.g., Deep Packet Inspection)
  - Intrusion Prevention (e.g., Firewalls)
  - Traffic Shaping/Quality of Service Enforcement
  - Load Balancing
  - Virtual Private Networks
  - Content Caching
  - Pre-network-access Authentication

Agenda

- ✓ Big data woes on the campus network
- Standard science DMZ solution
- Brief SDN overview
- A new DMZ approach
- Some results

How does one normally solve this problem?

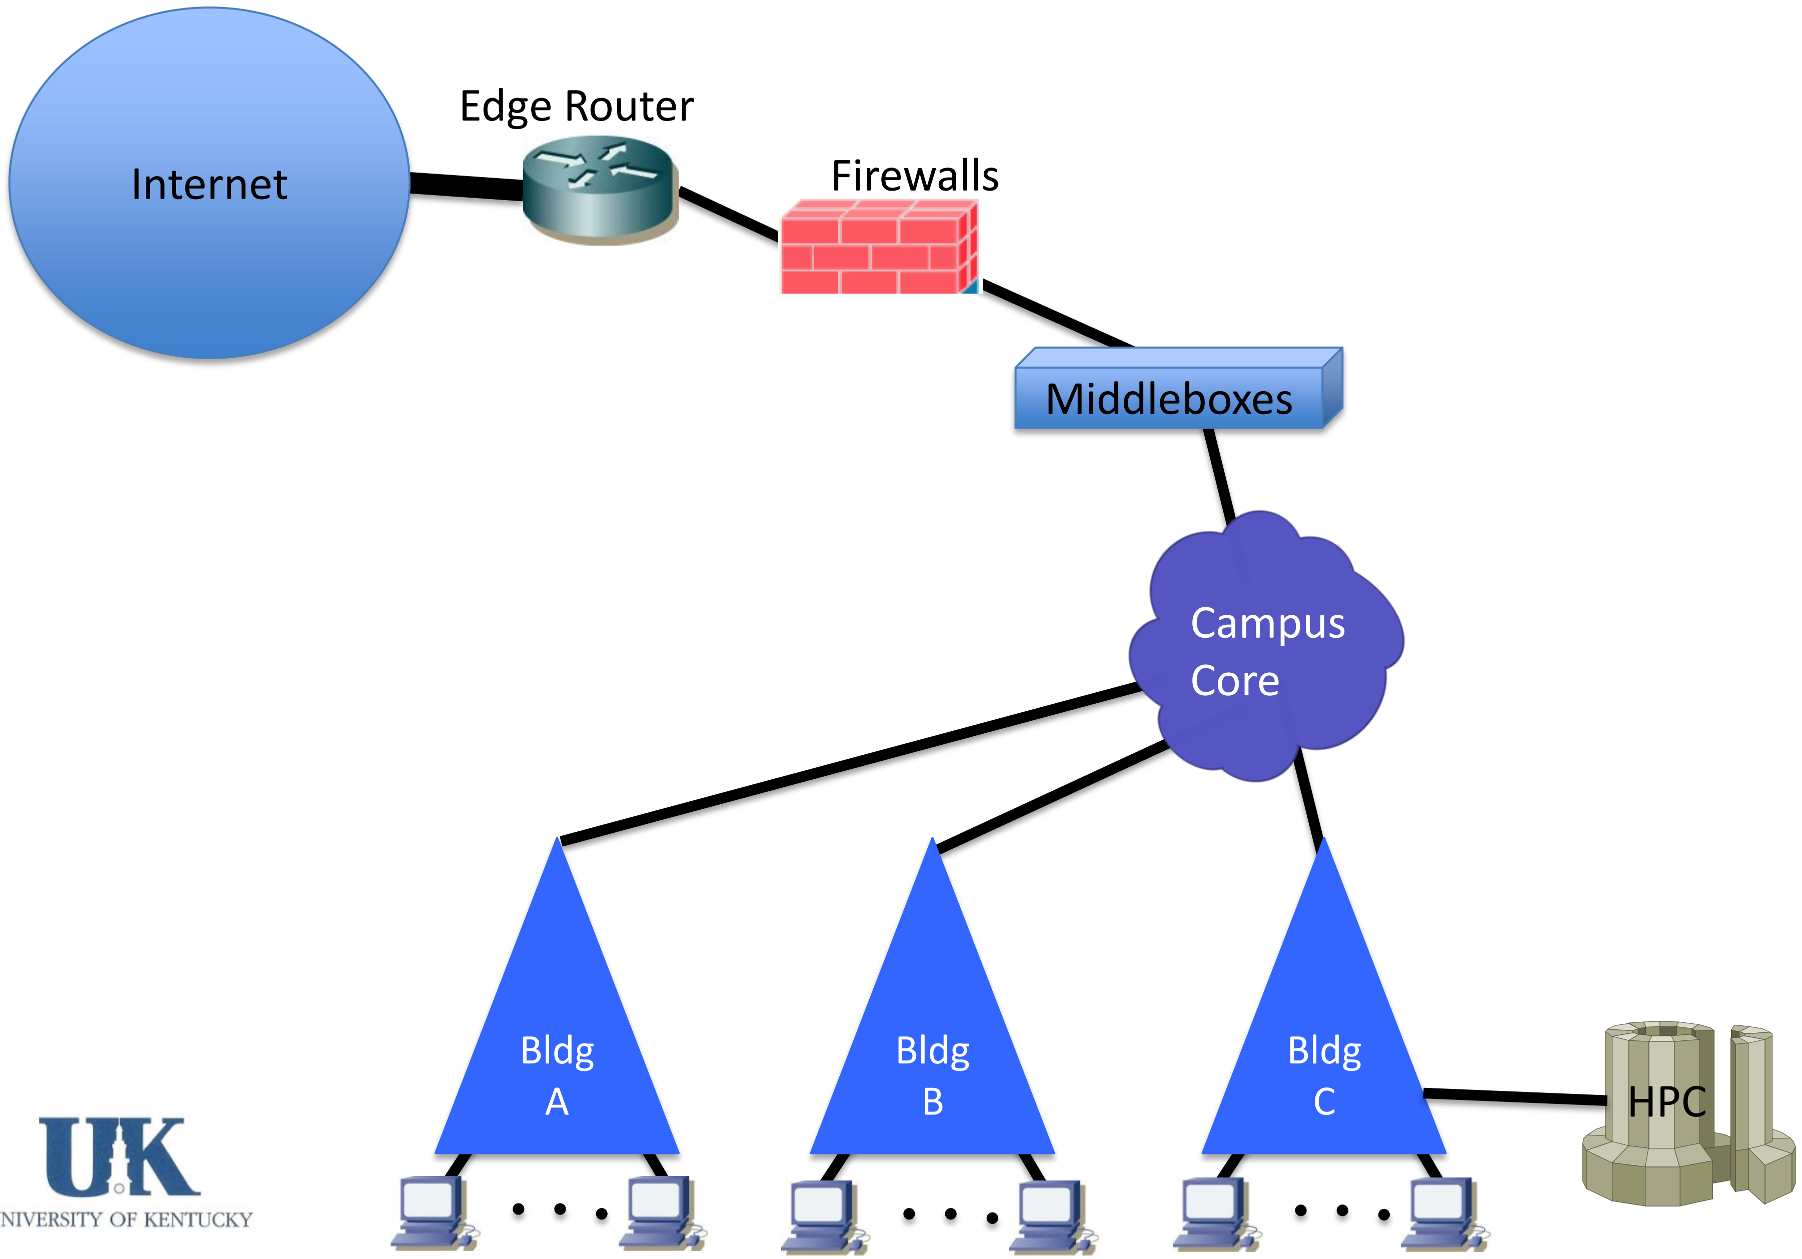

UNIVERSITY OF KENTUCKY

Move Nodes Outside the Firewall

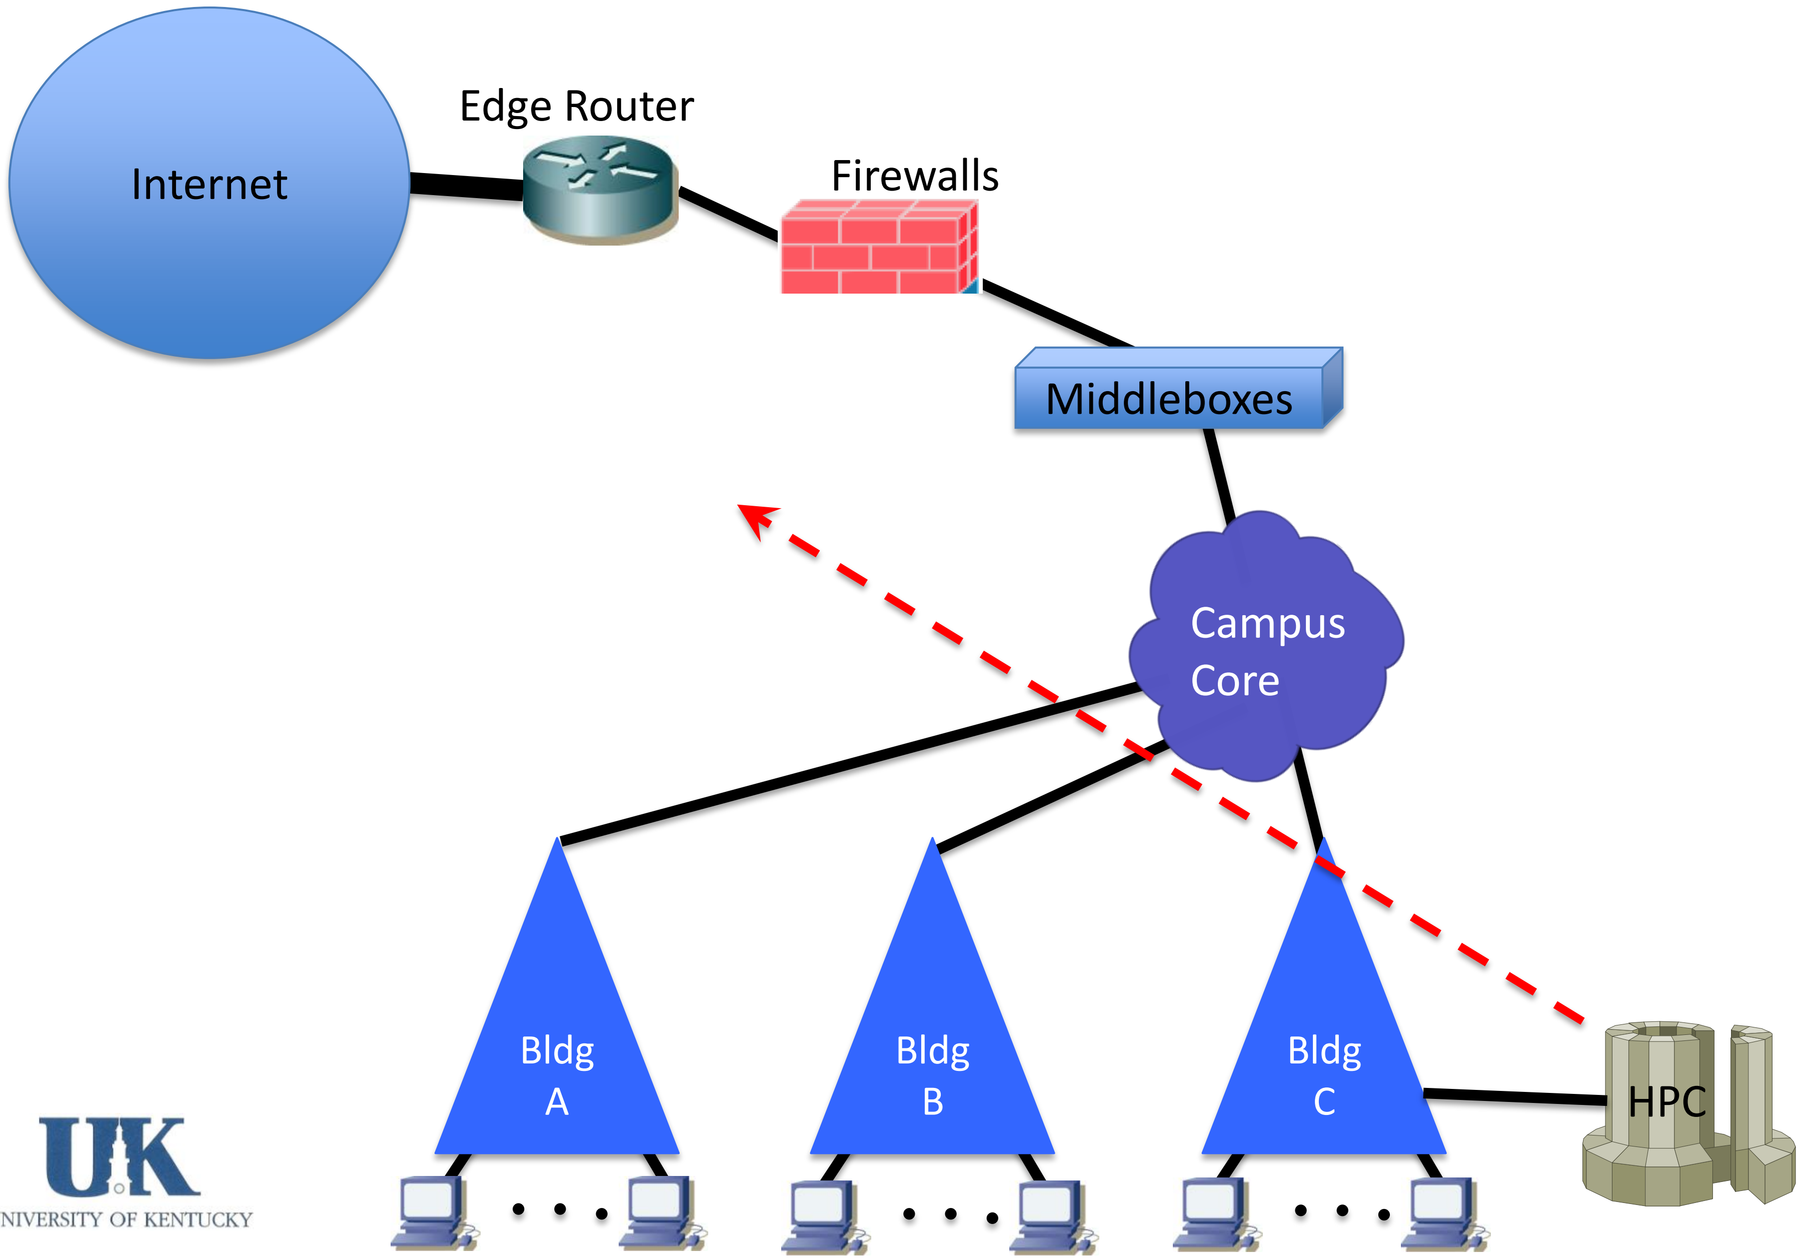

UK
UNIVERSITY OF KENTUCKY

Move Nodes Outside the Firewall

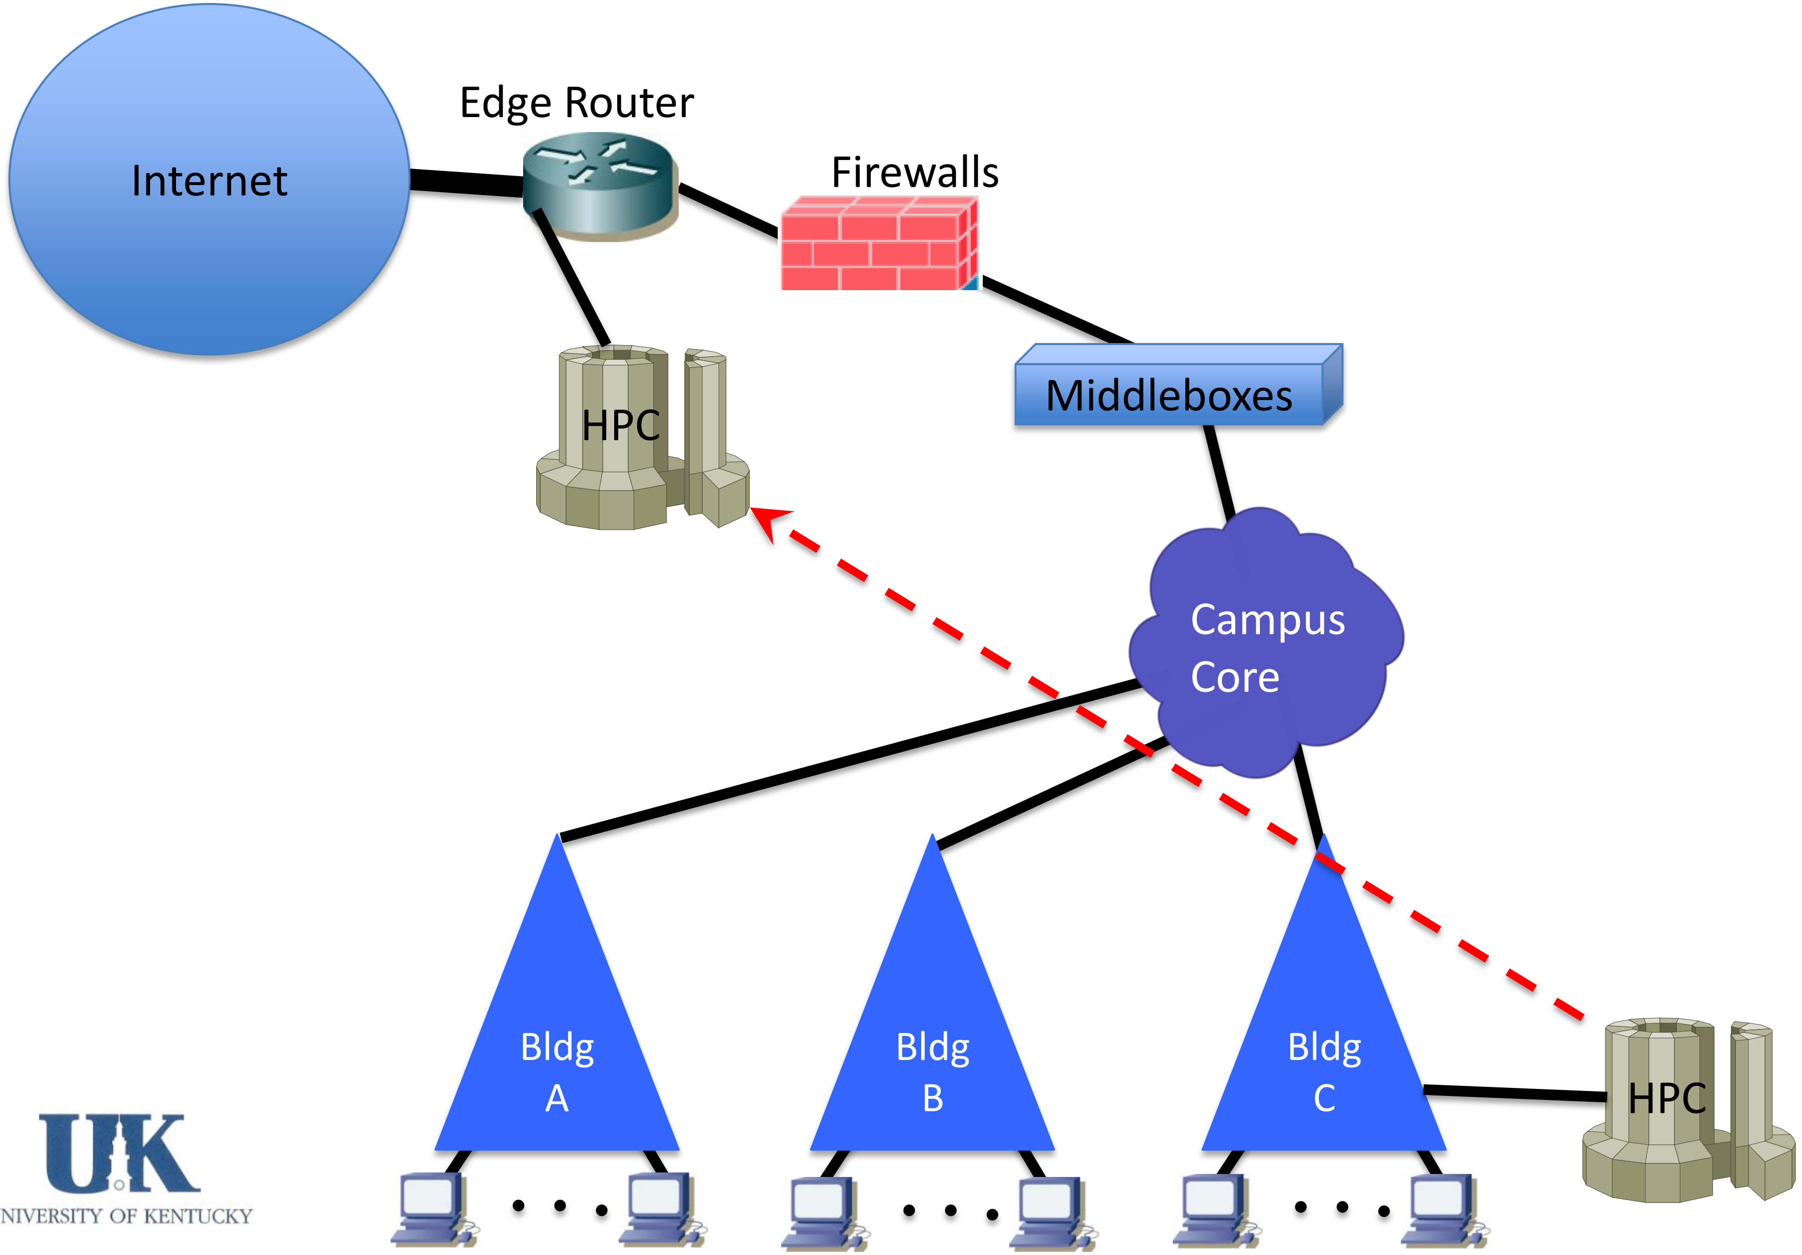

UNIVERSITY OF KENTUCKY

Move Nodes Outside the Firewall

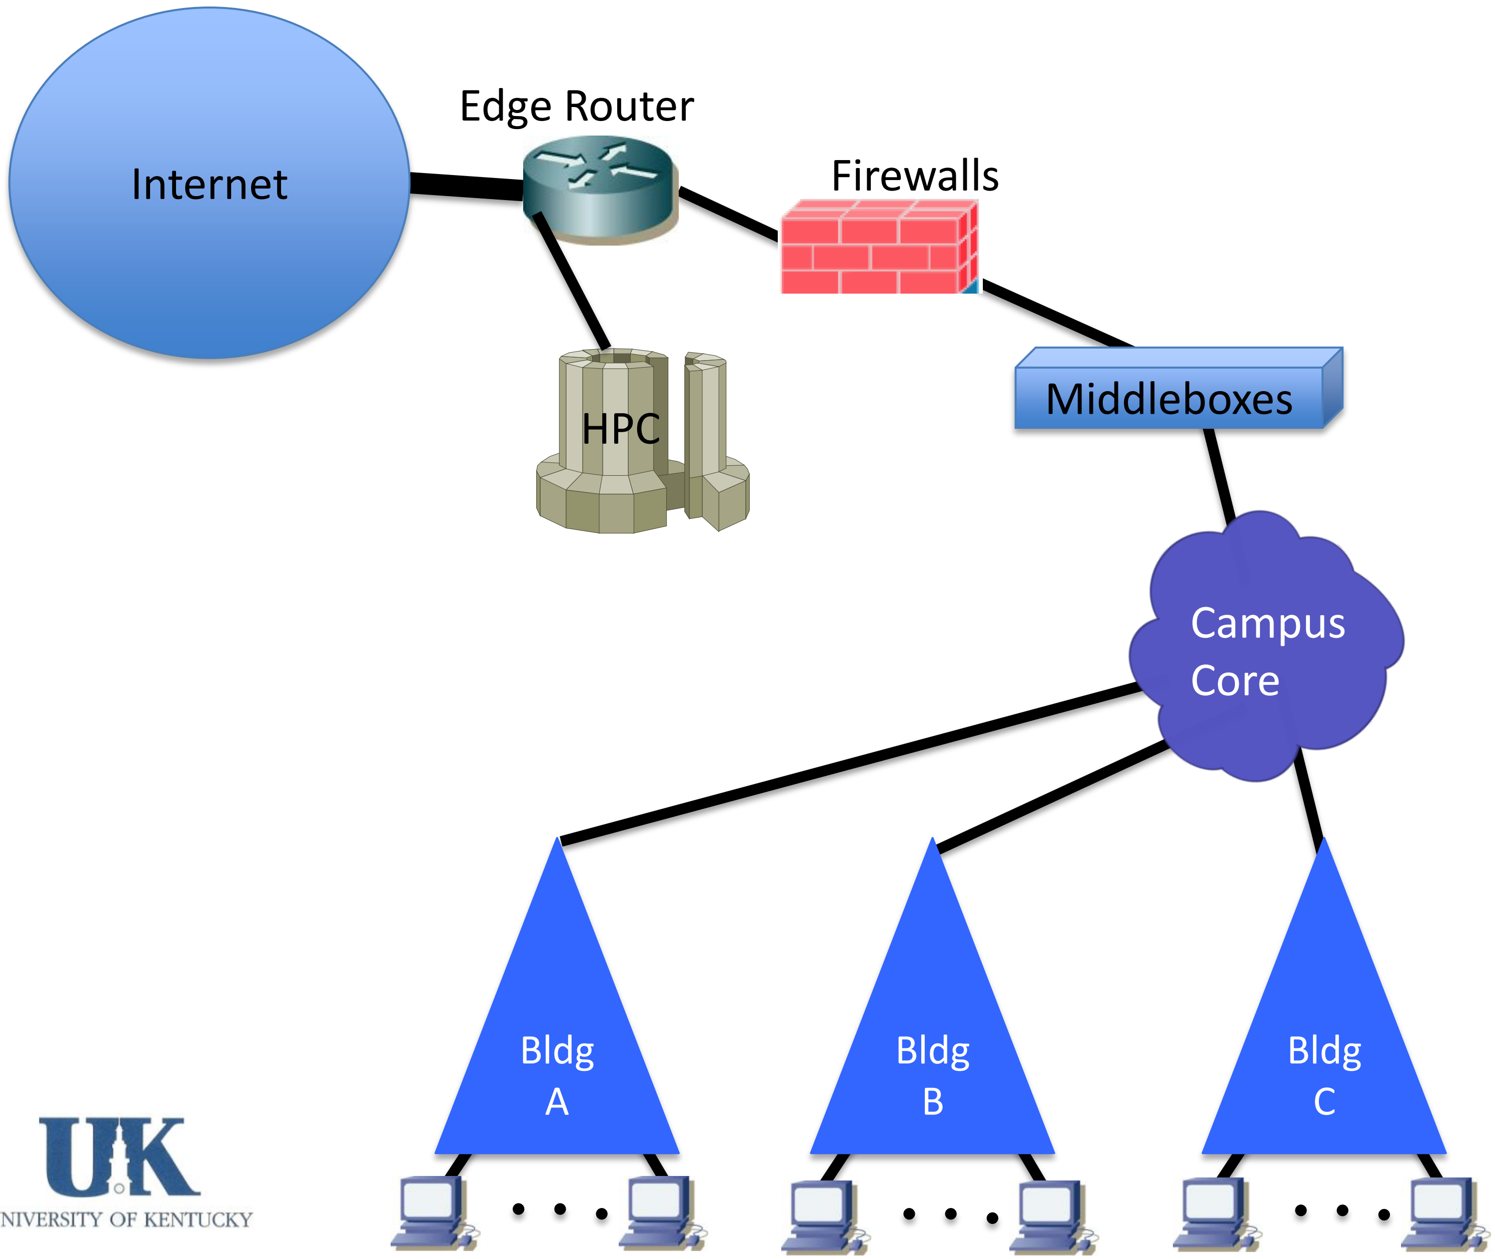

Campus Science DMZ

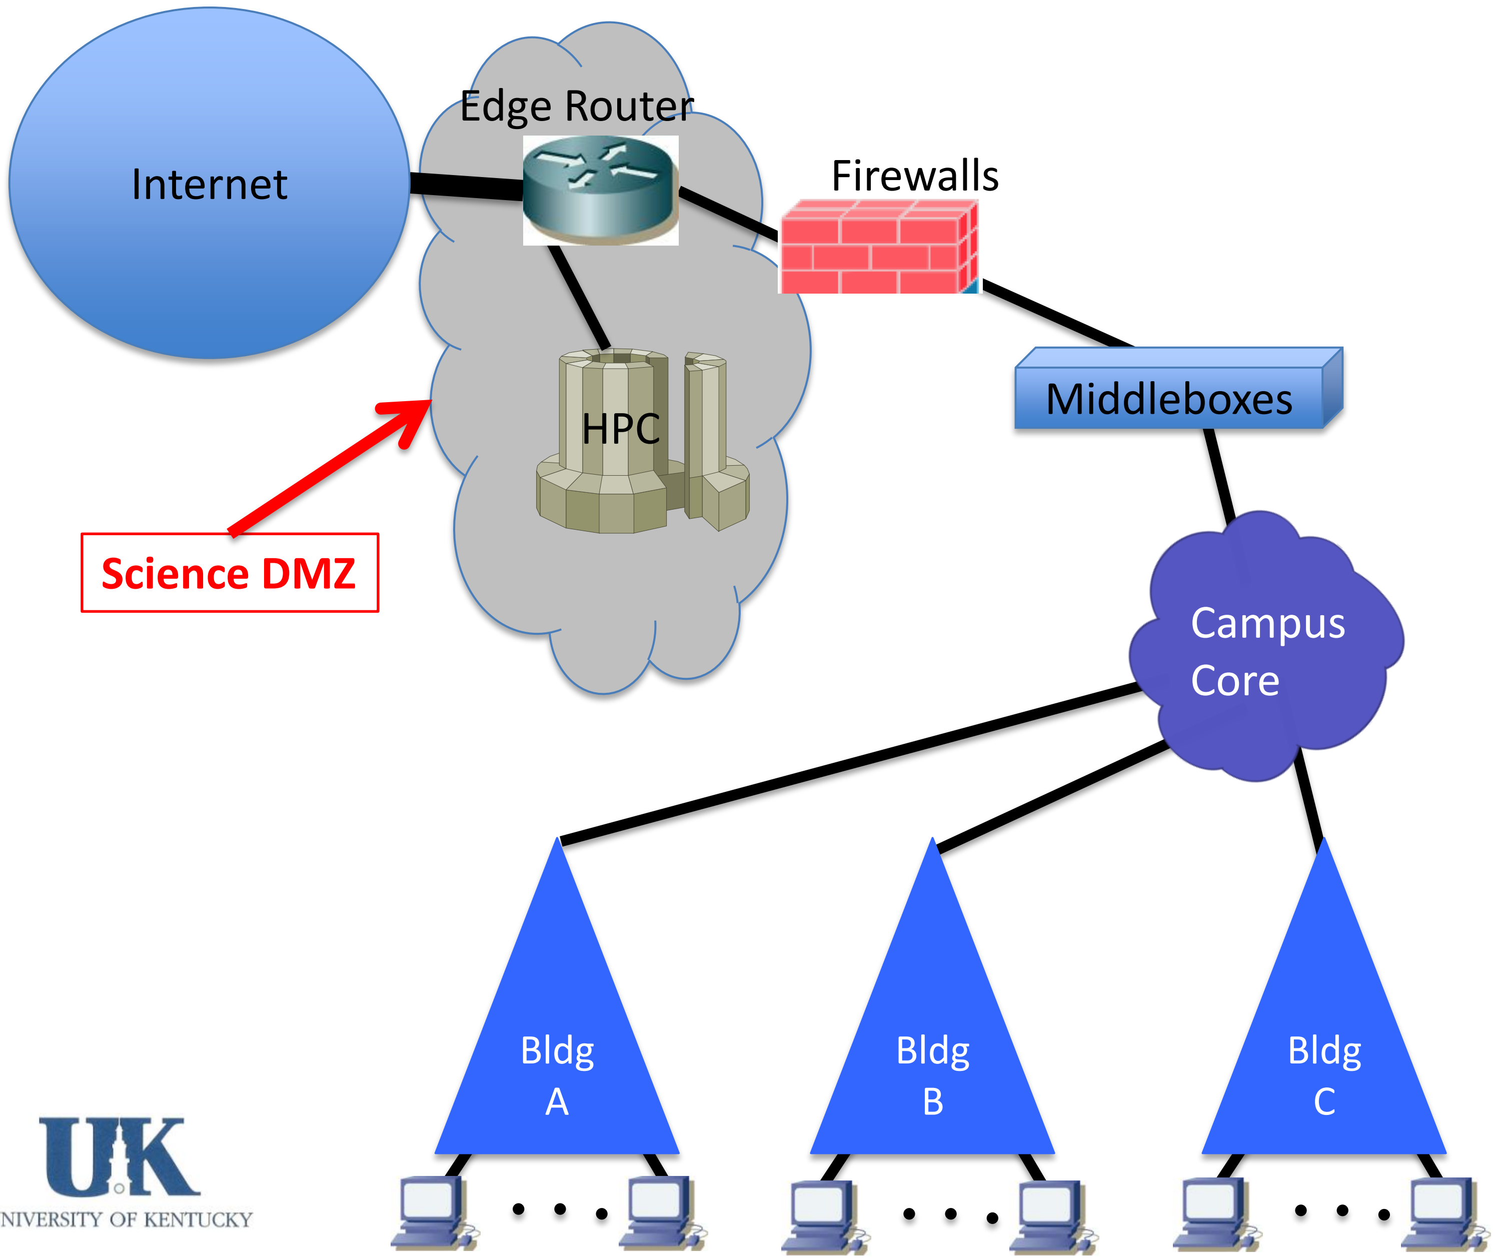

UK
UNIVERSITY OF KENTUCKY

Campus Science DMZ

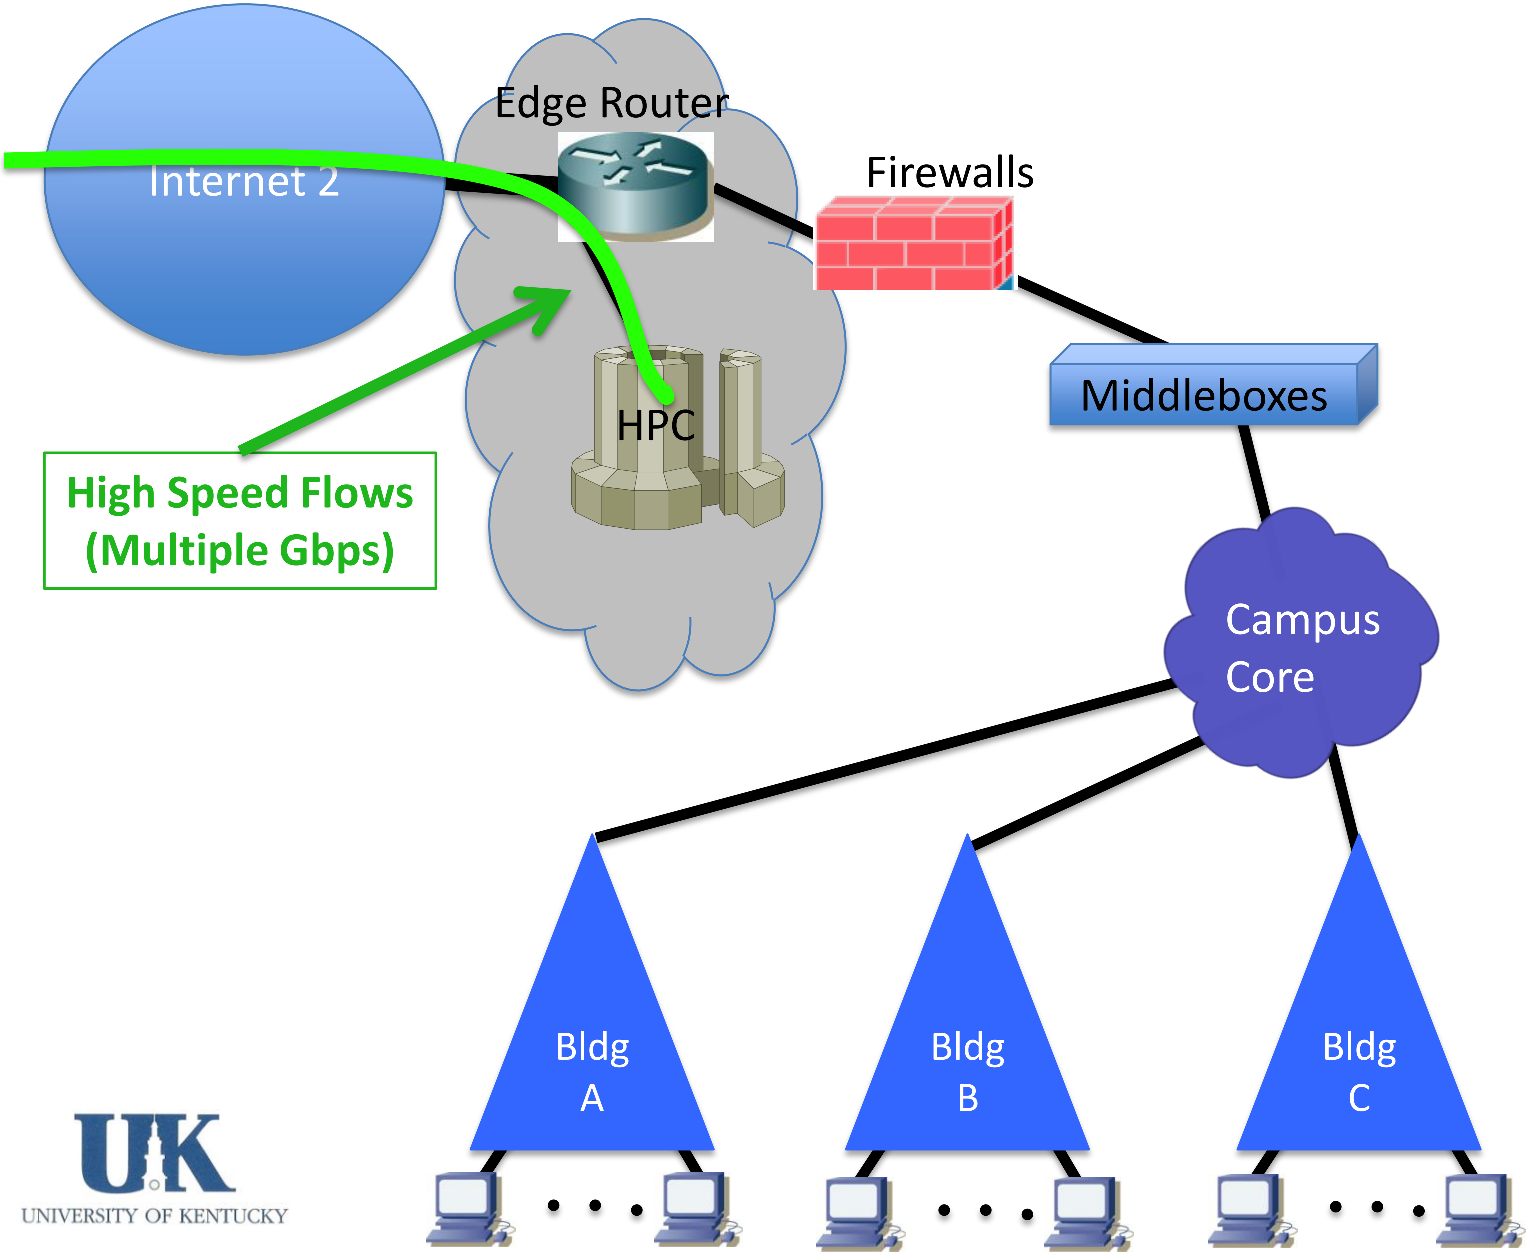

UNIVERSITY OF KENTUCKY

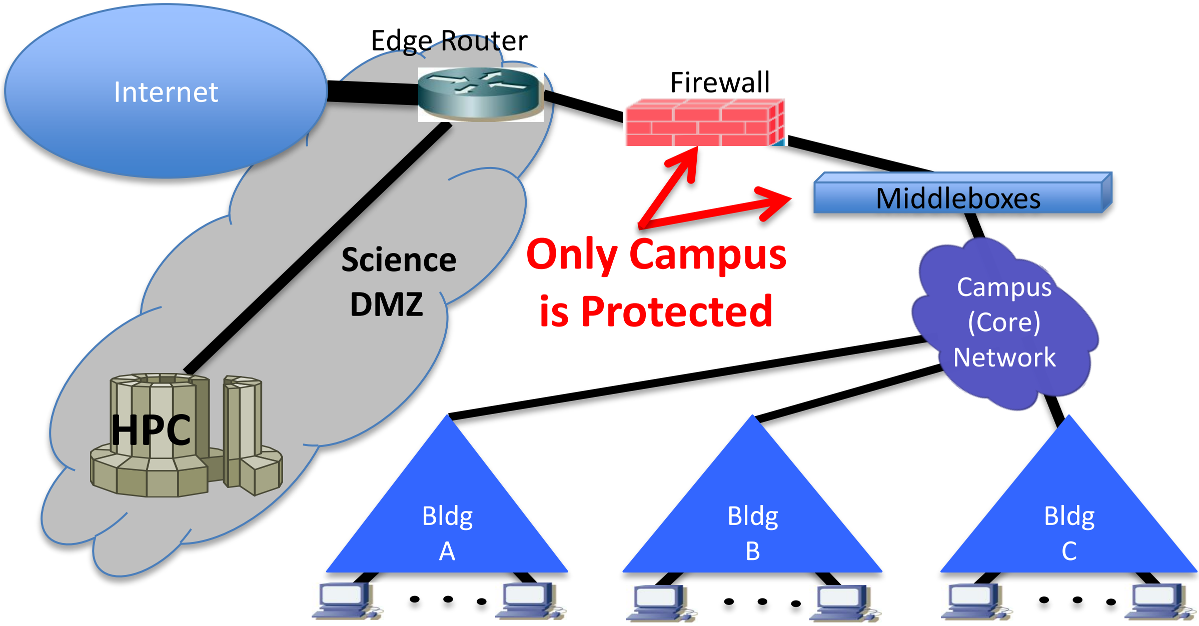

- Disadvantages:
  - Science DMZ machines are not protected by middleboxes.
  - Campus (middlebox) policy enforcement is not applied to any traffic from Science DMZ machines. Even non-science flows (e.g., Netflix) bypass campus policy enforcement.
  - Researchers must decide whether to connect their machines to the Science DMZ or the Campus Network.

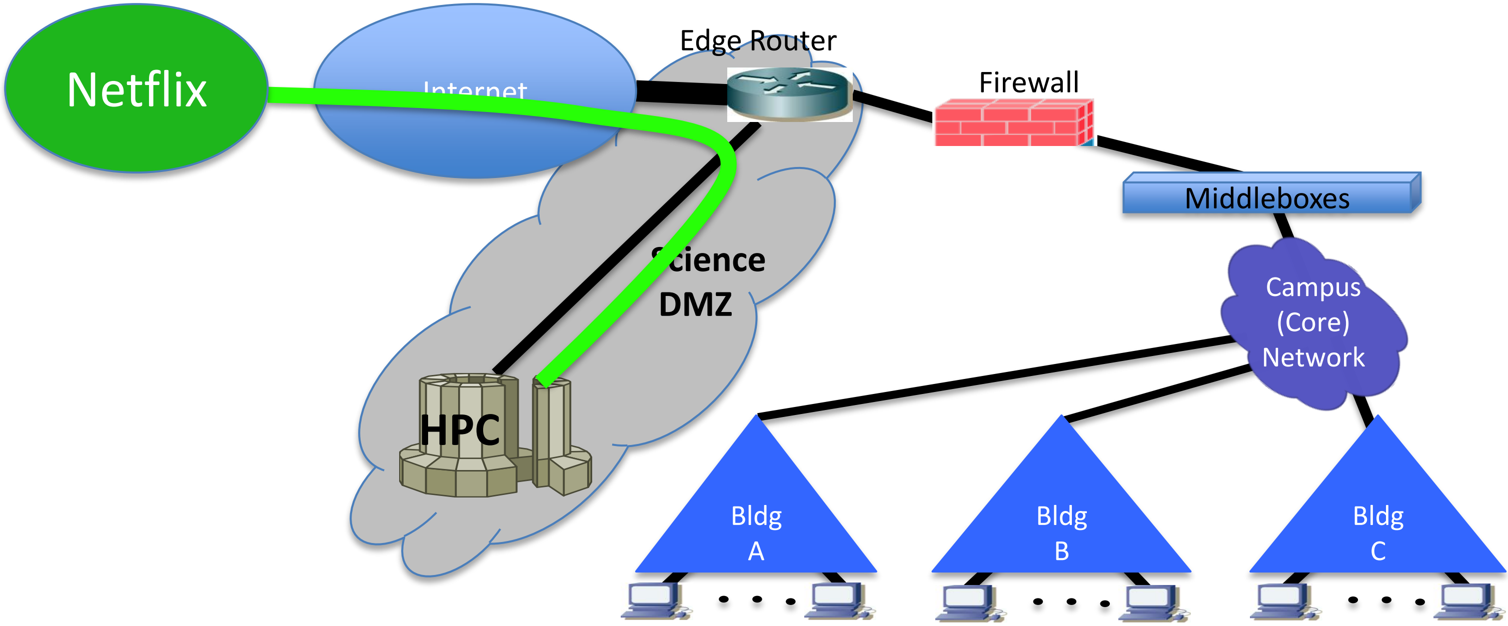

- Disadvantages:
  - Science DMZ machines are not protected by middleboxes.
  - Campus (middlebox) policy enforcement is not applied to any traffic from Science DMZ machines. Even non-science flows (e.g., Netflix) bypass campus policy enforcement.
  - Researchers must decide whether to connect their machines to the Science DMZ or the Campus Network.

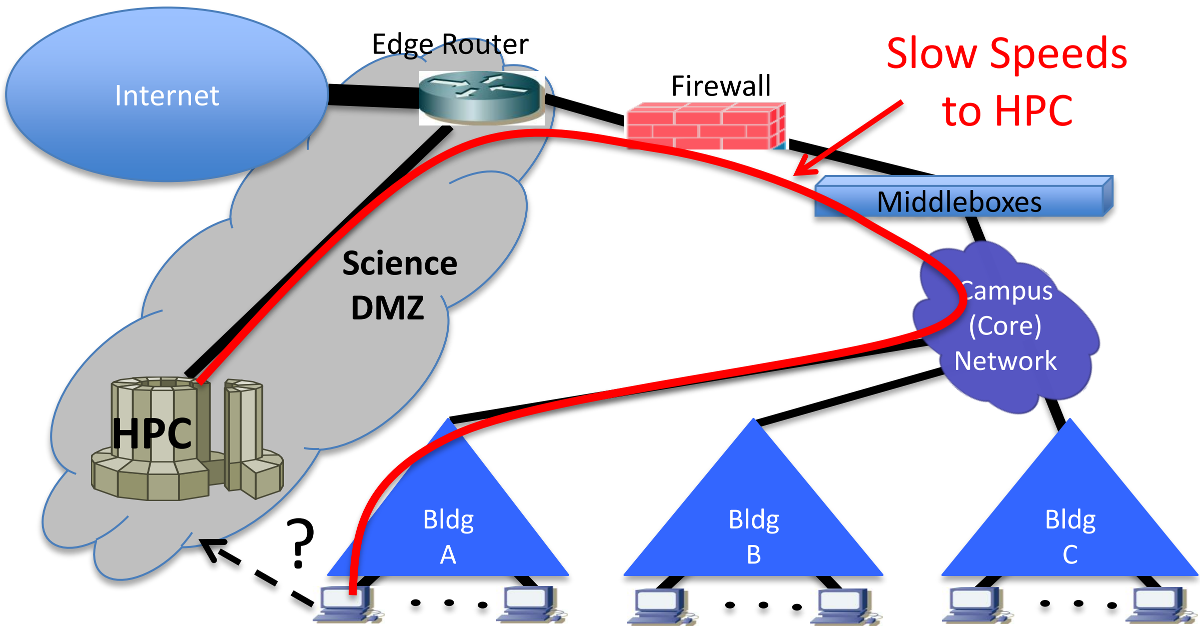

- Disadvantages:
  - Science DMZ machines are not protected by middleboxes.
  - Campus (middlebox) policy enforcement is not applied to any traffic from Science DMZ machines. Even non-science flows (e.g., Netflix) bypass campus policy enforcement.
  - Researchers must decide whether to connect their machines to the Science DMZ or the Campus Network.

Agenda

✓ Big data woes on the campus network

✓ Standard science DMZ solution

- Brief SDN overview
- A new DMZ approach
- Some results

UNIVERSITY OF KENTUCKY

Normal Switch

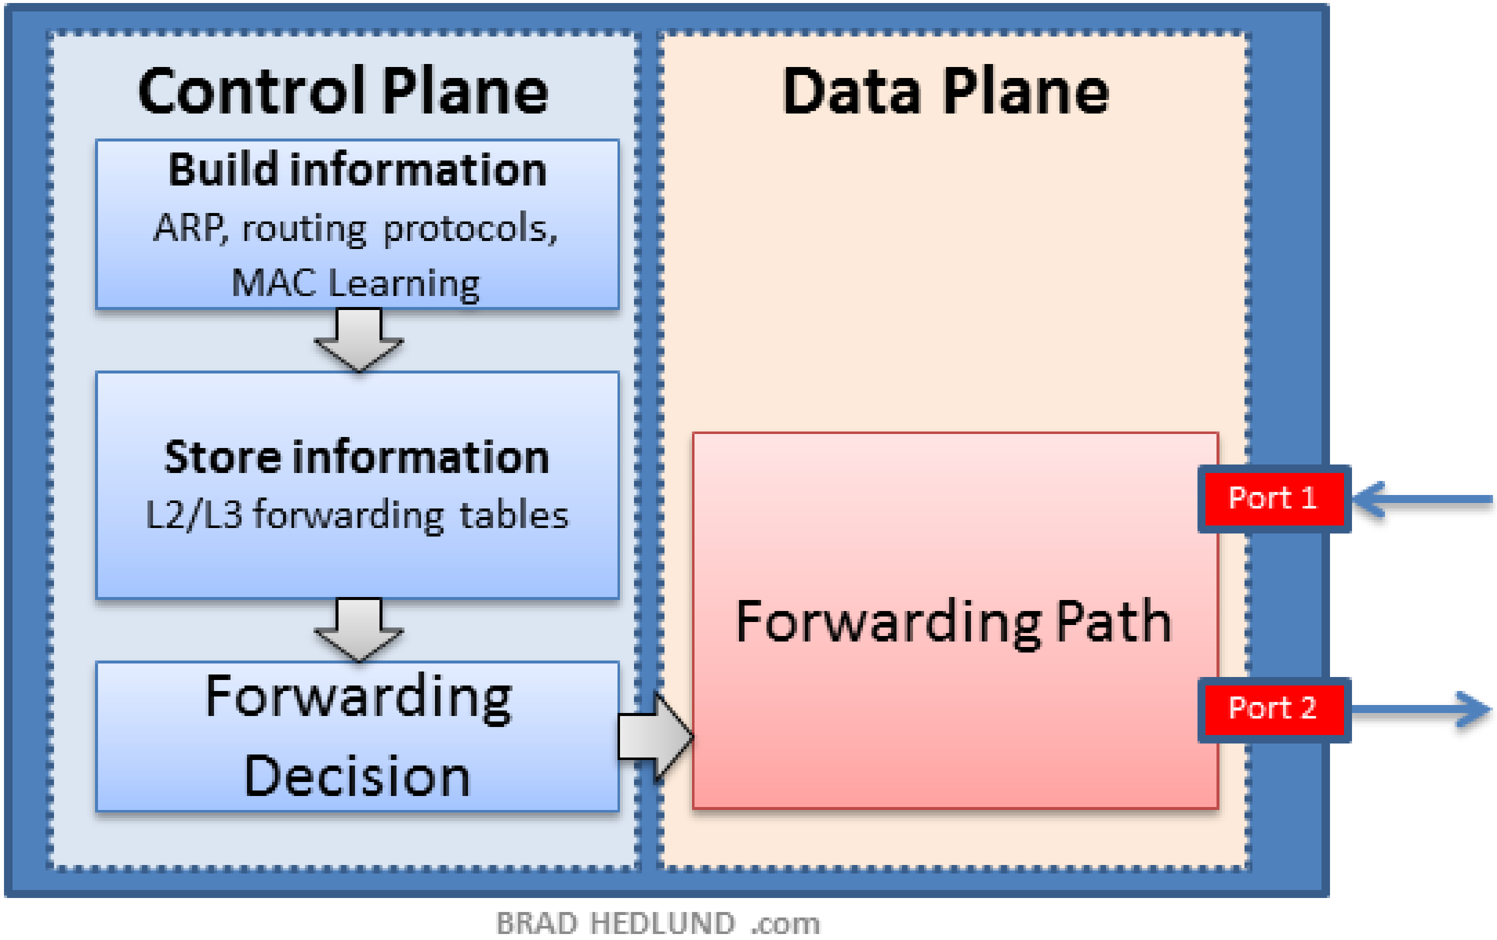

UNIVERSITY OF KENTUCKY

OpenFlow Enabled Switch

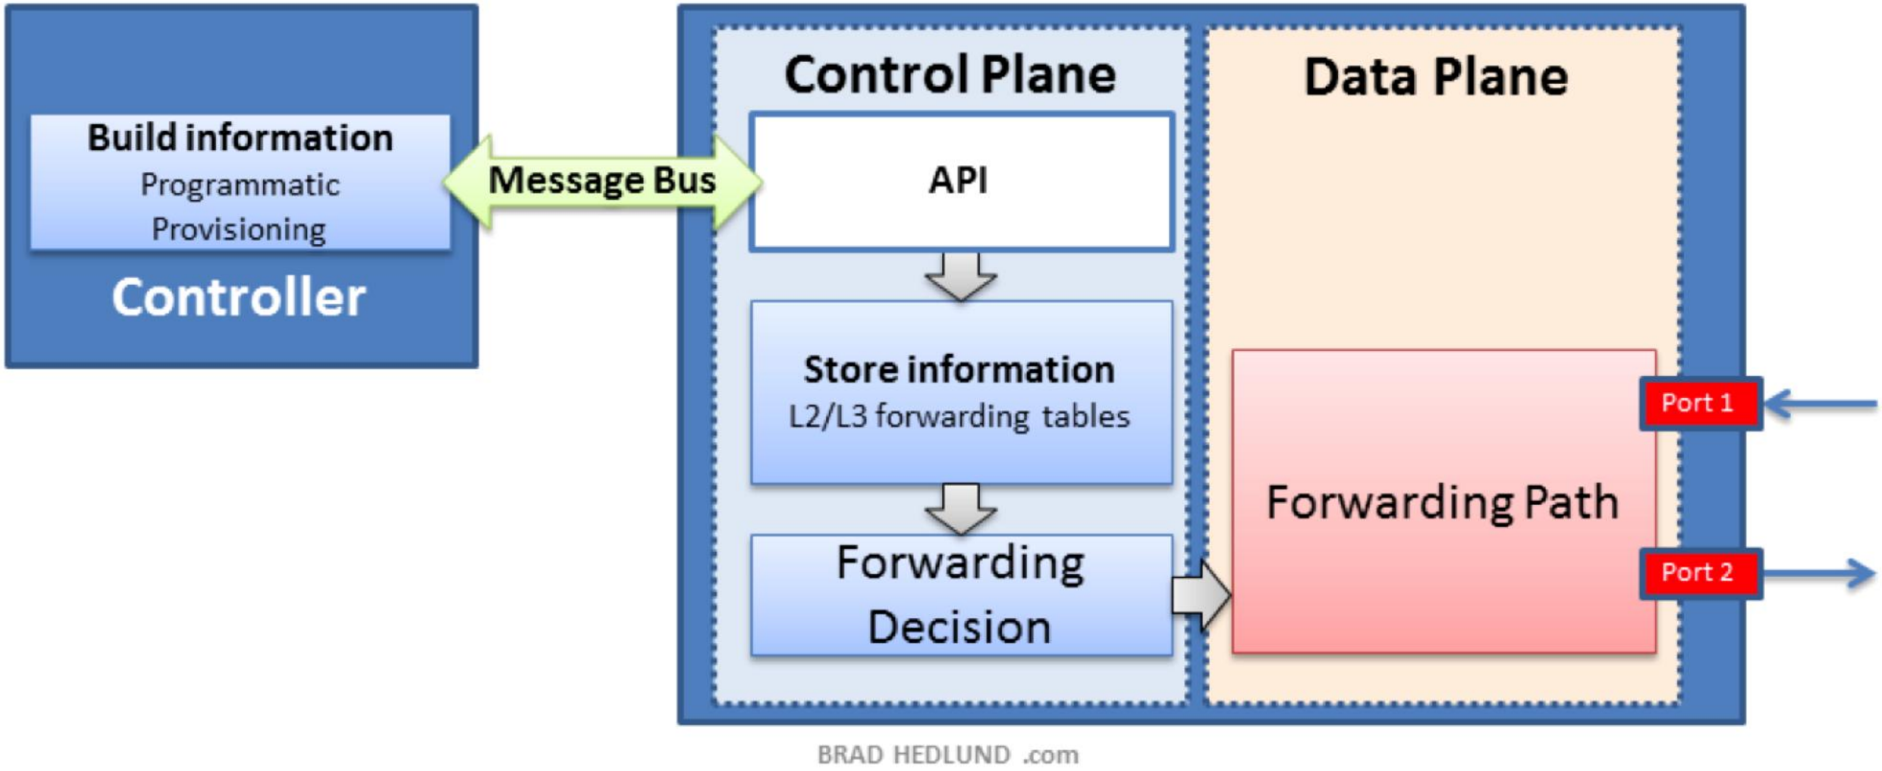

OpenFlow : physical separation of network control plane from the data plane

UNIVERSITY OF KENTUCKY

OpenFlow : physical separation of network control plane from the data plane

UNIVERSITY OF KENTUCKY

OpenFlow Rules : Match Packet then Take Actions

Match:

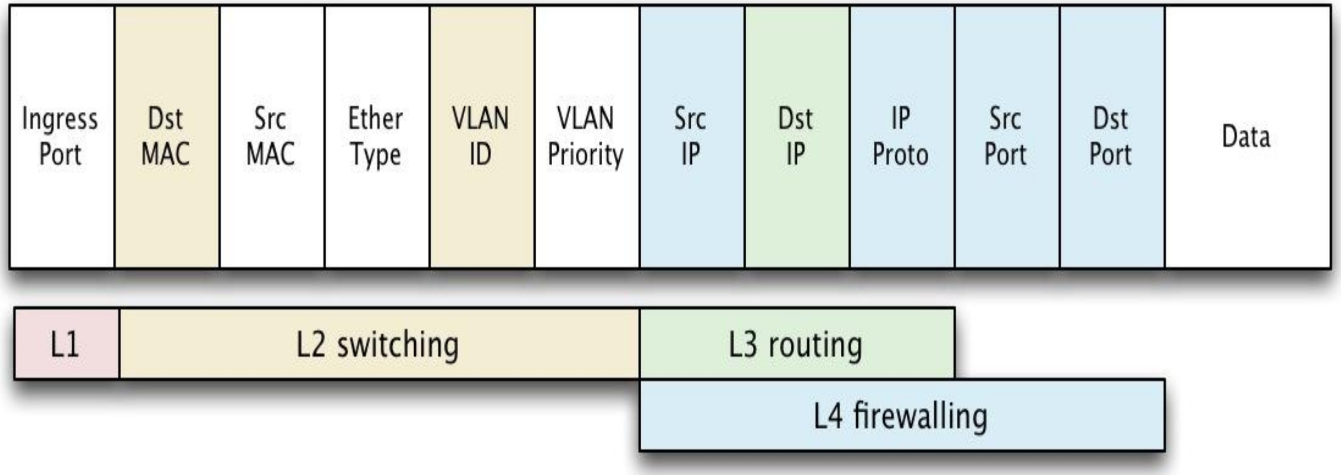

Ingress Port	Dst MAC	Src MAC	Ether Type	VLAN ID	VLAN Priority	Src IP	Dst IP	IP Proto	Src Port	Dst Port	Data

L1

L2 switching

L3 routing

L4 firewalling

Actions:

Drop

Forward – to port, flood, to controller, normal…

Set – mac, vlan id, ip address…

How Does It Work?

Controller pushes initial OpenFlow rules to every switch

Match	Action
bddp	forward to controller (topology discovery)
dhcp	forward to controller & Normal (end node discovery)
arp	forward to controller & Normal (end node discovery)
*	Normal

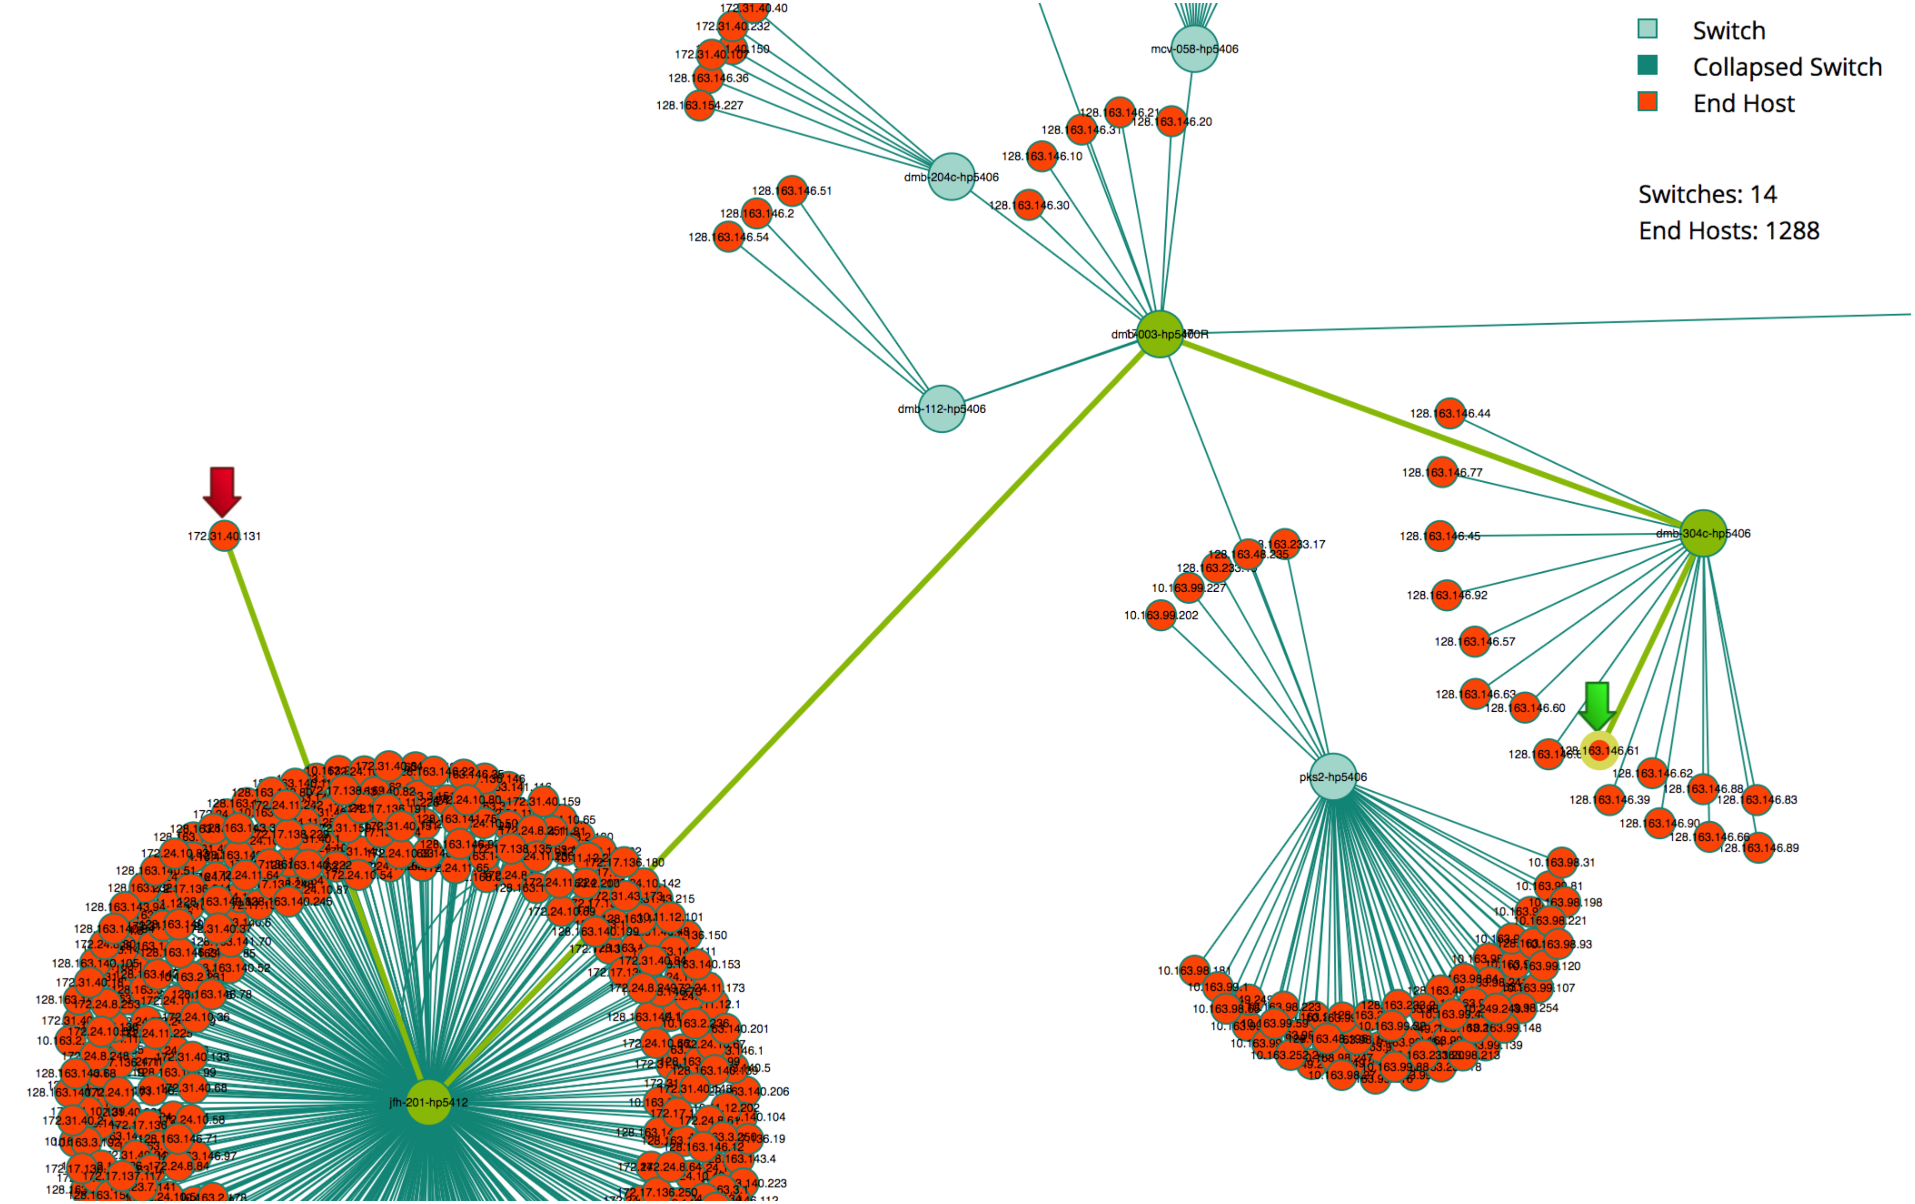
🟩 Switch
🟩 Collapsed Switch
🟥 End Host
Switches: 14
End Hosts: 1288
mcv-058-hp5406
dmb-204c-hp5406
dmb-003-hp5406R
dmb-112-hp5406
dmb-304c-hp5406
pks2-hp5406
jfh-201-hp5412
172.31.40.40
172.31.40.232
172.31.40.107
128.163.146.36
128.163.154.227
128.163.146.51
128.163.146.2
128.163.146.54
128.163.146.21
128.163.146.31
128.163.146.20
128.163.146.10
128.163.146.30
128.163.146.44
128.163.146.77
128.163.146.45
128.163.146.92
128.163.146.57
128.163.146.63
128.163.146.60
128.163.146.61
128.163.146.62
128.163.146.88
128.163.146.83
128.163.146.39
128.163.146.90
128.163.146.66
128.163.146.89
172.31.40.131
128.163.233.17
128.163.48.235
10.163.99.227
10.163.99.202
10.163.98.31

University of Kentucky SDN Implementation

- Long term goal: all campus networking equipment will be SDN capable
- Use Normal rule for most traffic
- Modify “special” traffic with OF when needed

- Examples of “special” traffic:
- Push large physics research data directly to Internet2
- Avoid middle box bottlenecks for special traffic
- Drop hostile traffic

UK
UNIVERSITY OF KENTUCKY

Agenda

- ✓ Big data woes on the campus network
- ✓ Standard science DMZ solution
- ✓ Brief SDN overview
- A new DMZ approach
- Some results

Developing an All-Campus Science DMZ

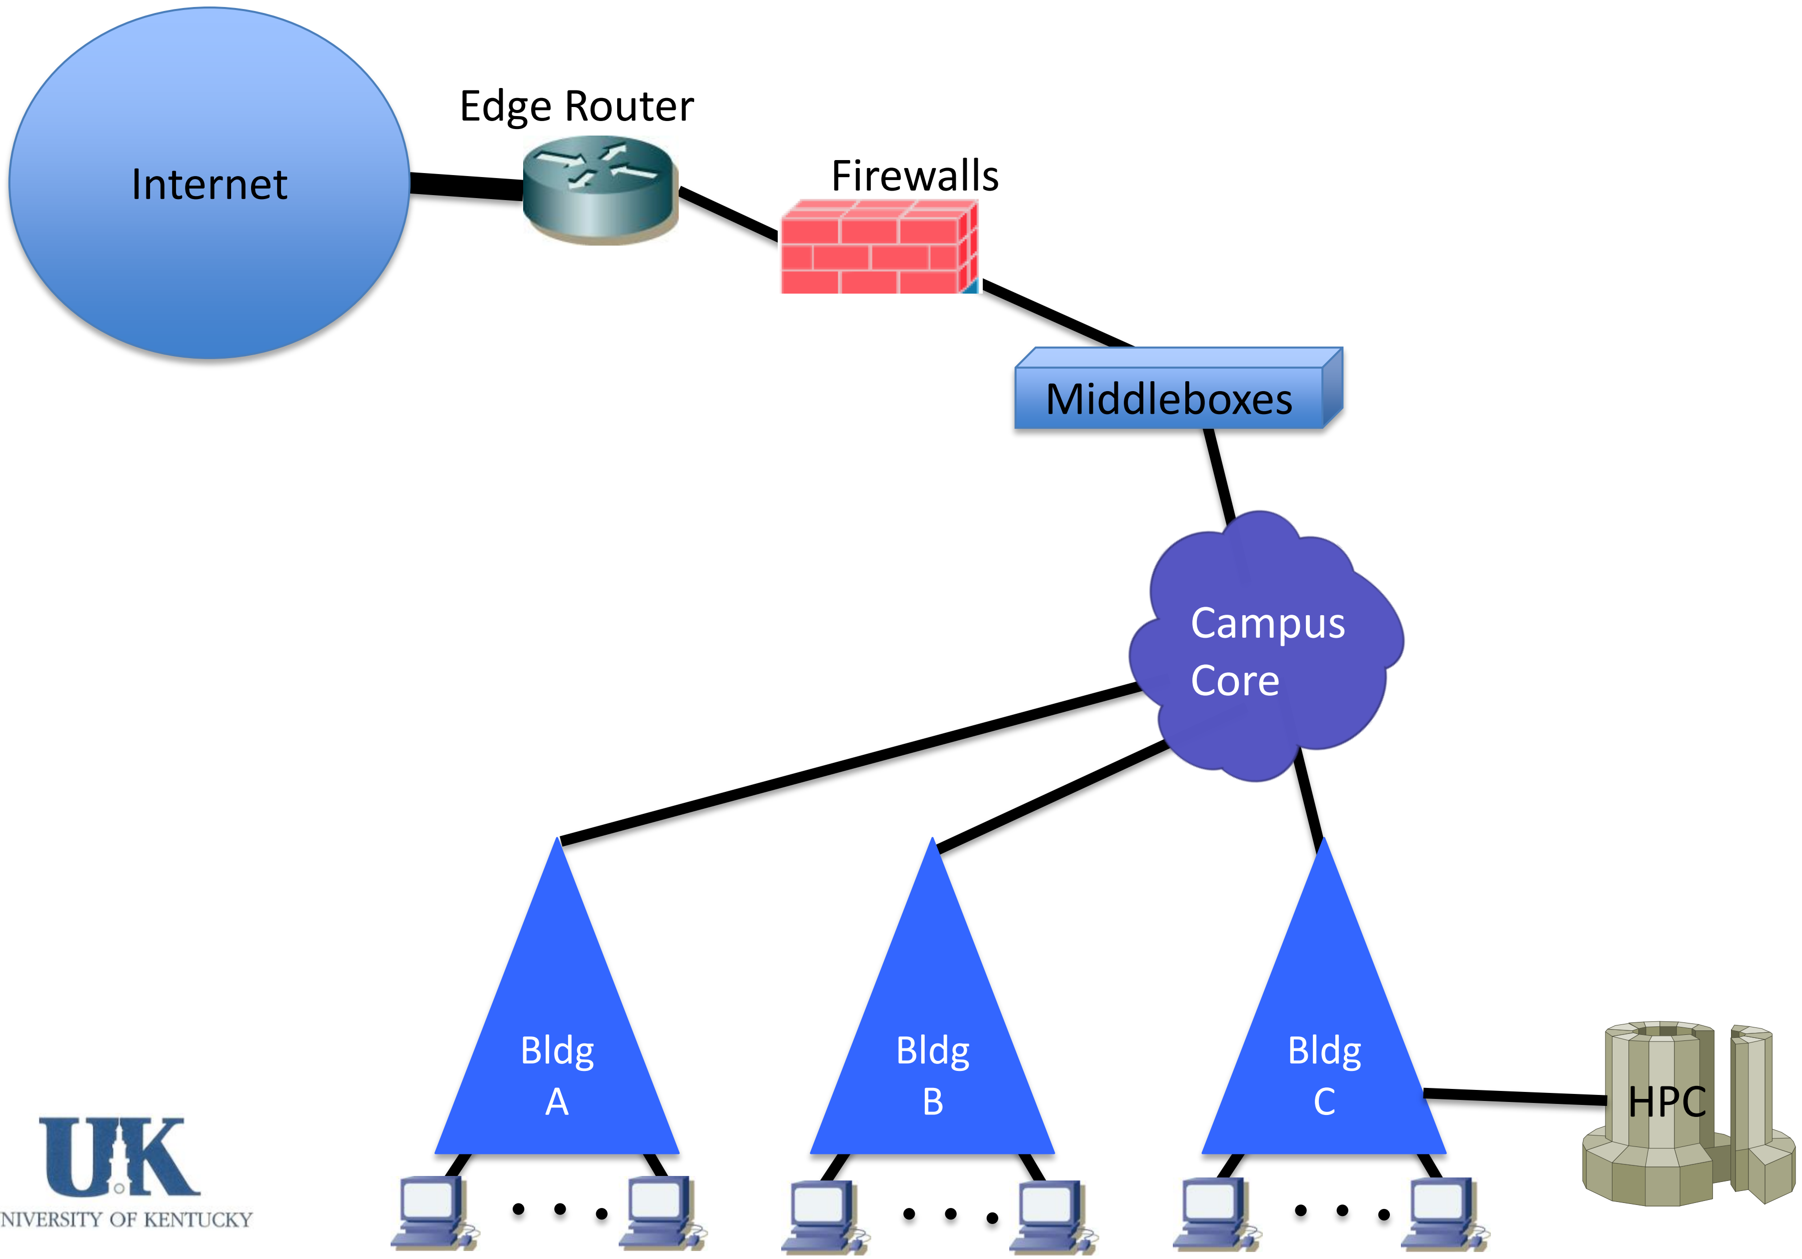

UNIVERSITY OF KENTUCKY

Developing an All-Campus Science DMZ

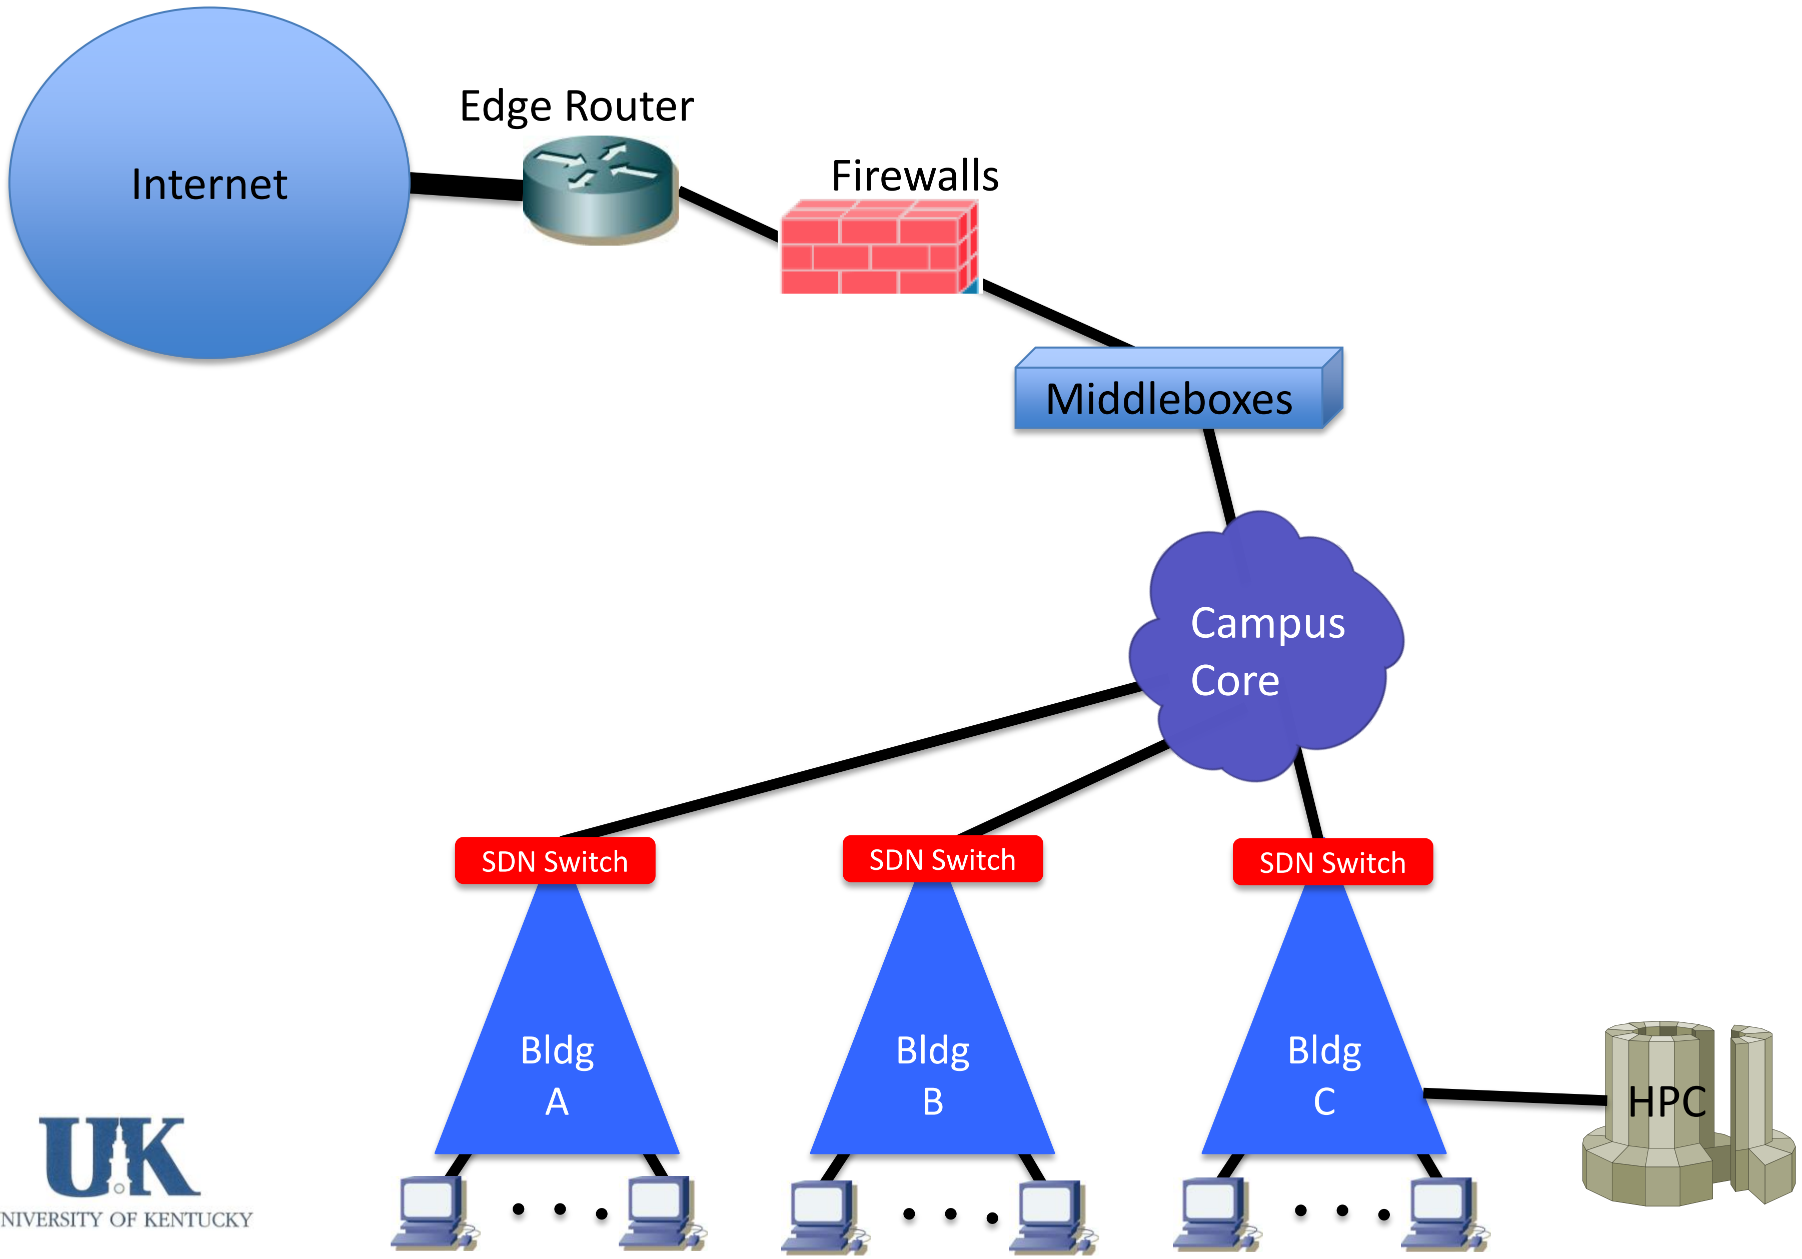

UNIVERSITY OF KENTUCKY

Developing an All-Campus Science DMZ

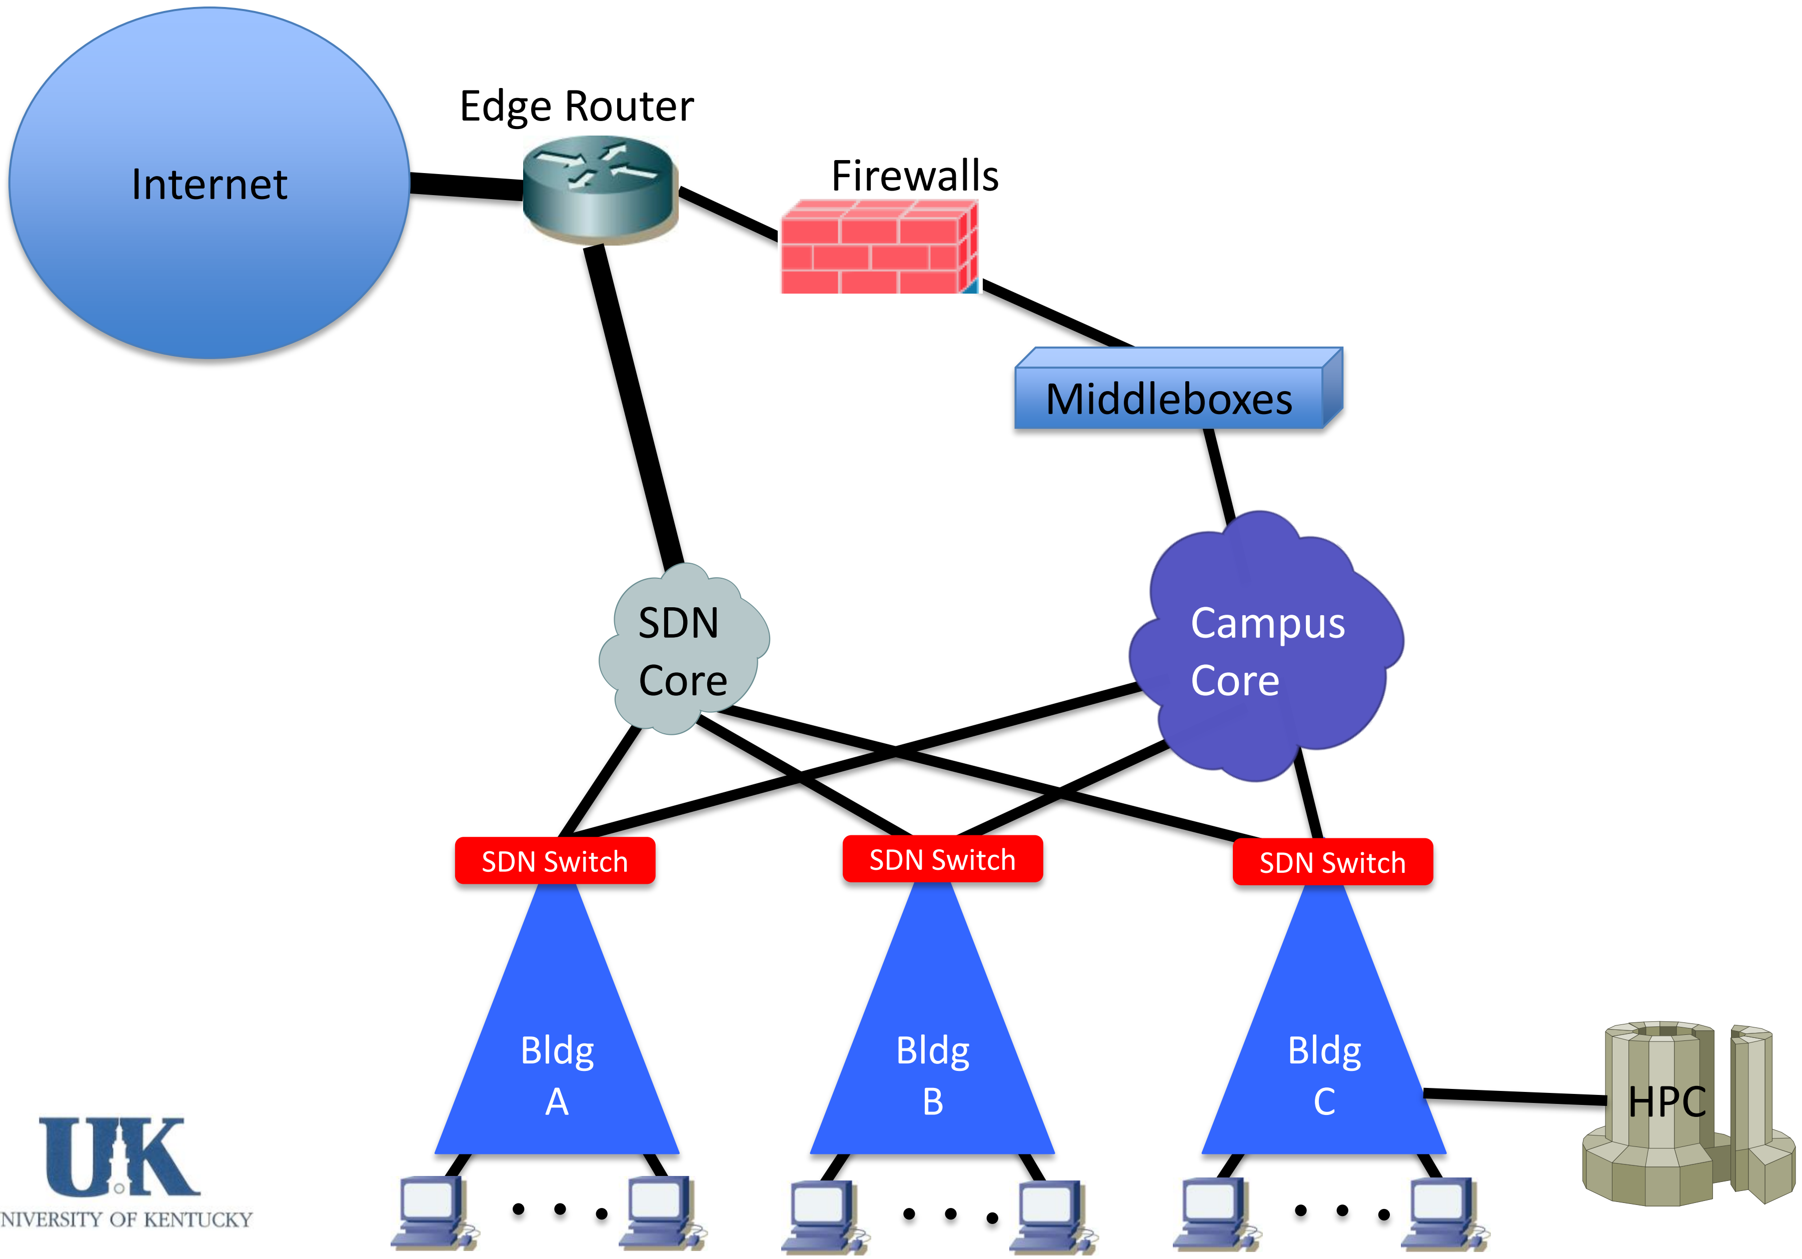

UNIVERSITY OF KENTUCKY

Developing an All-Campus Science DMZ

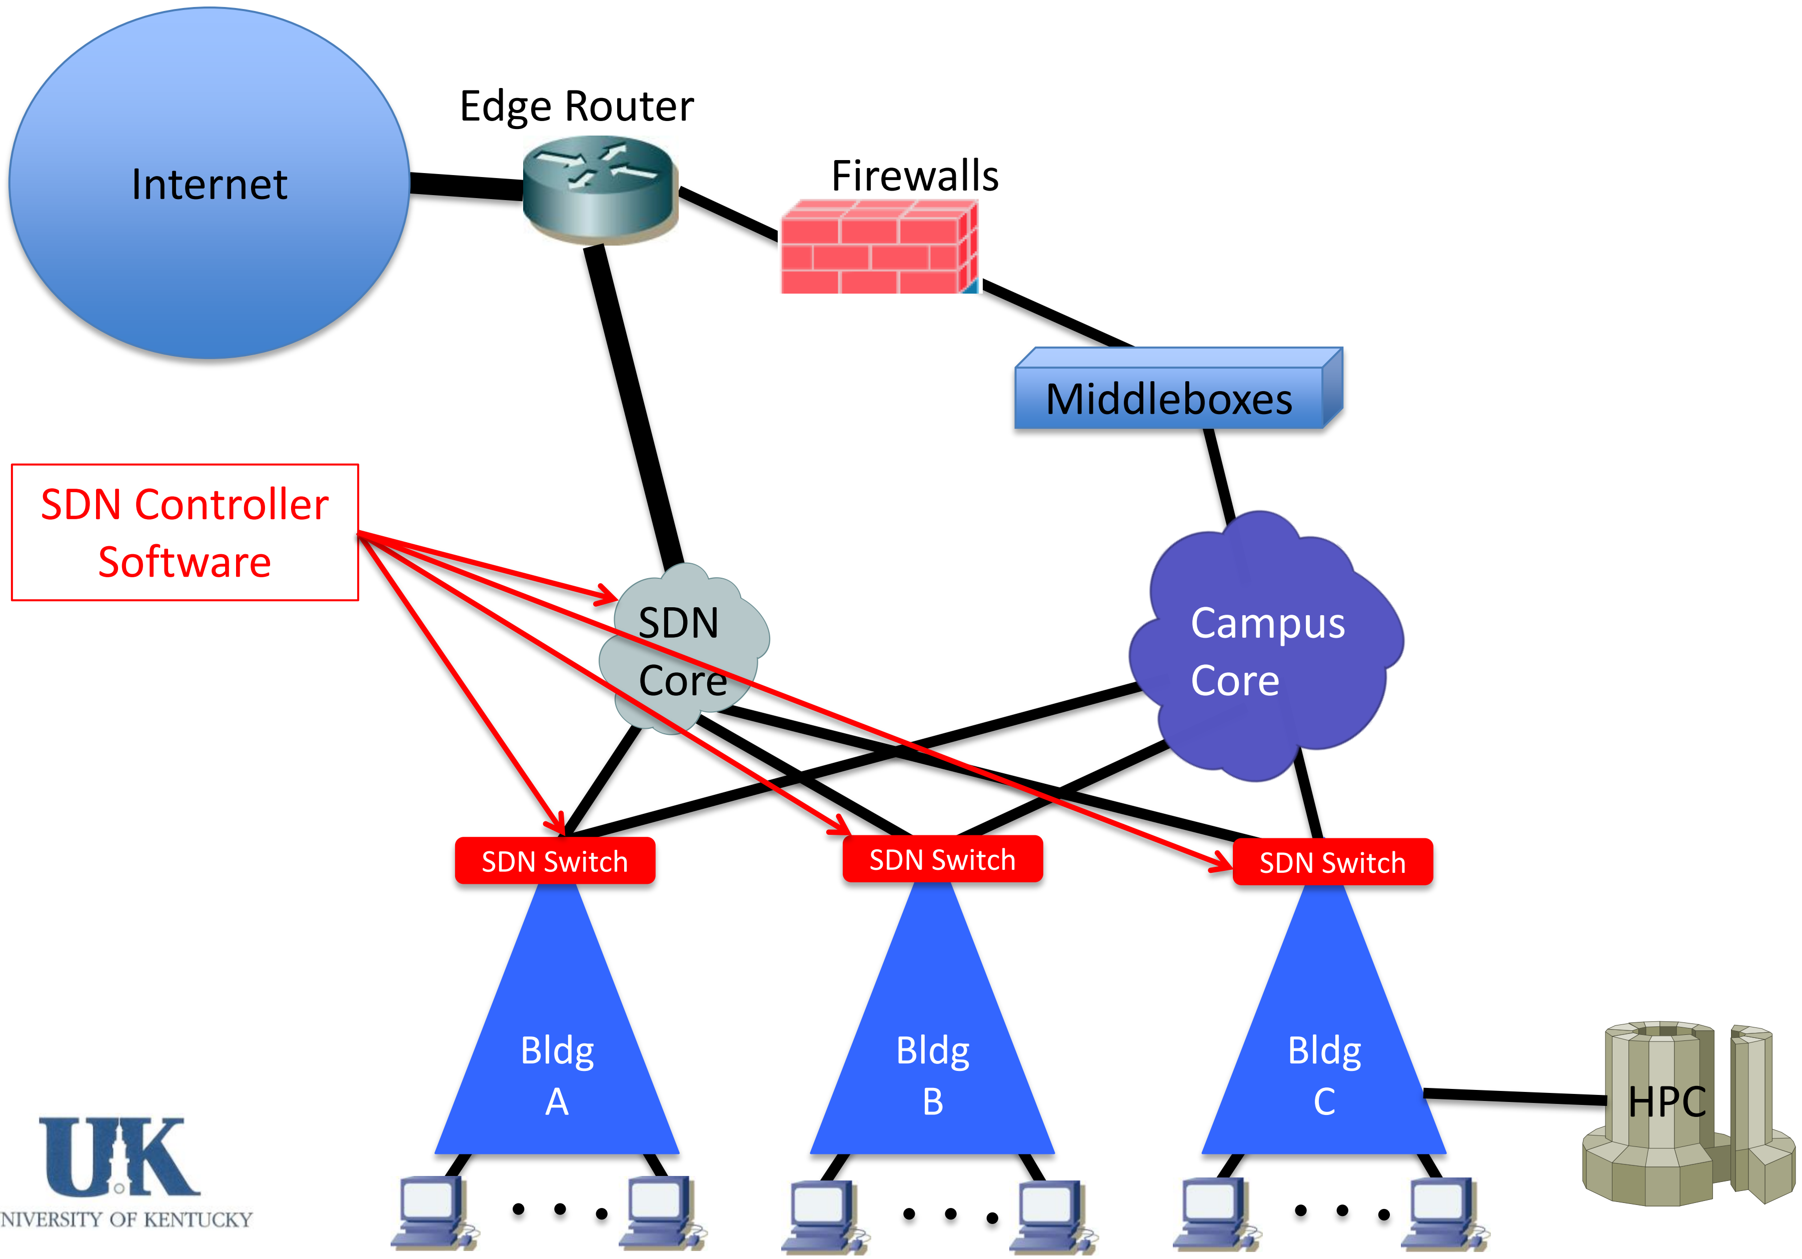

UNIVERSITY OF KENTUCKY

Developing an All-Campus Science DMZ

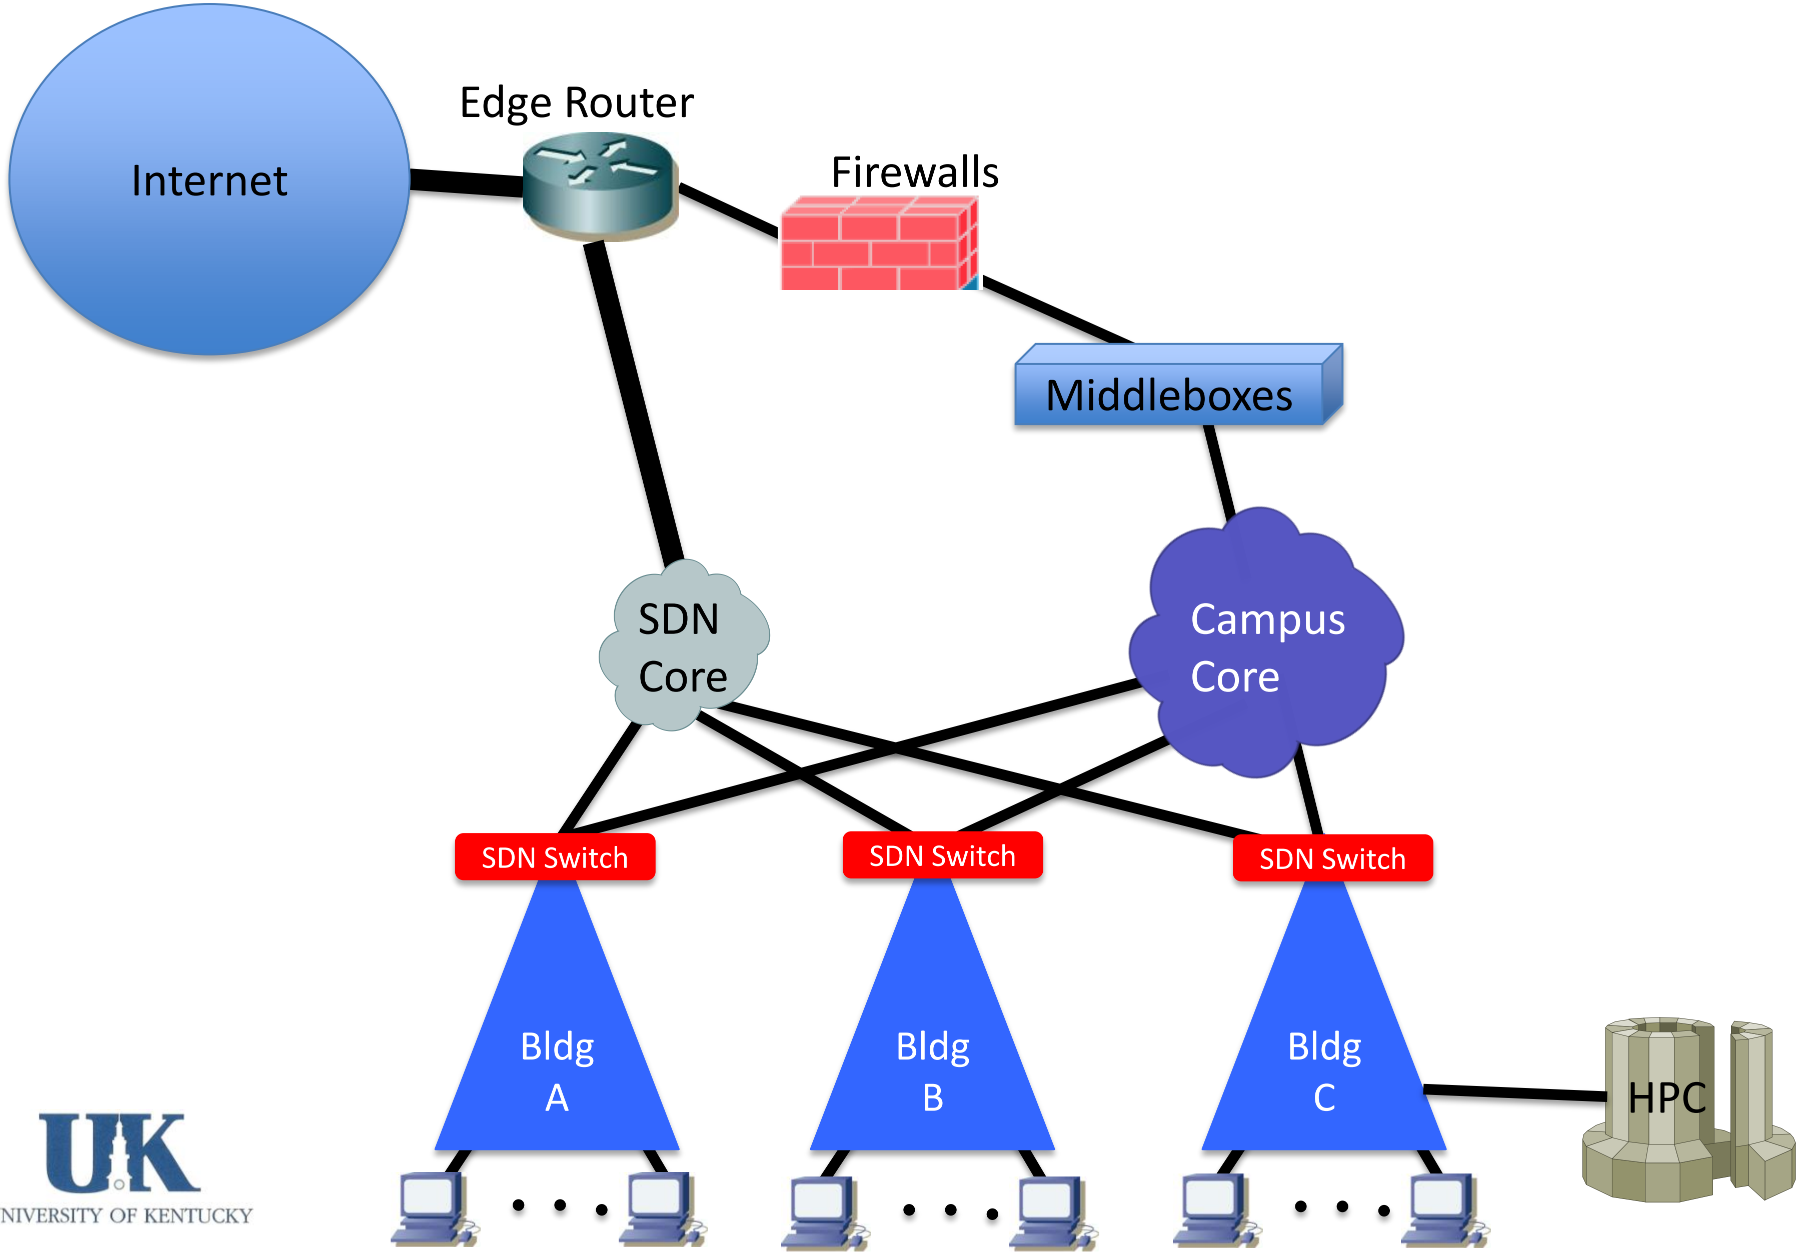

Developing an All-Campus Science DMZ

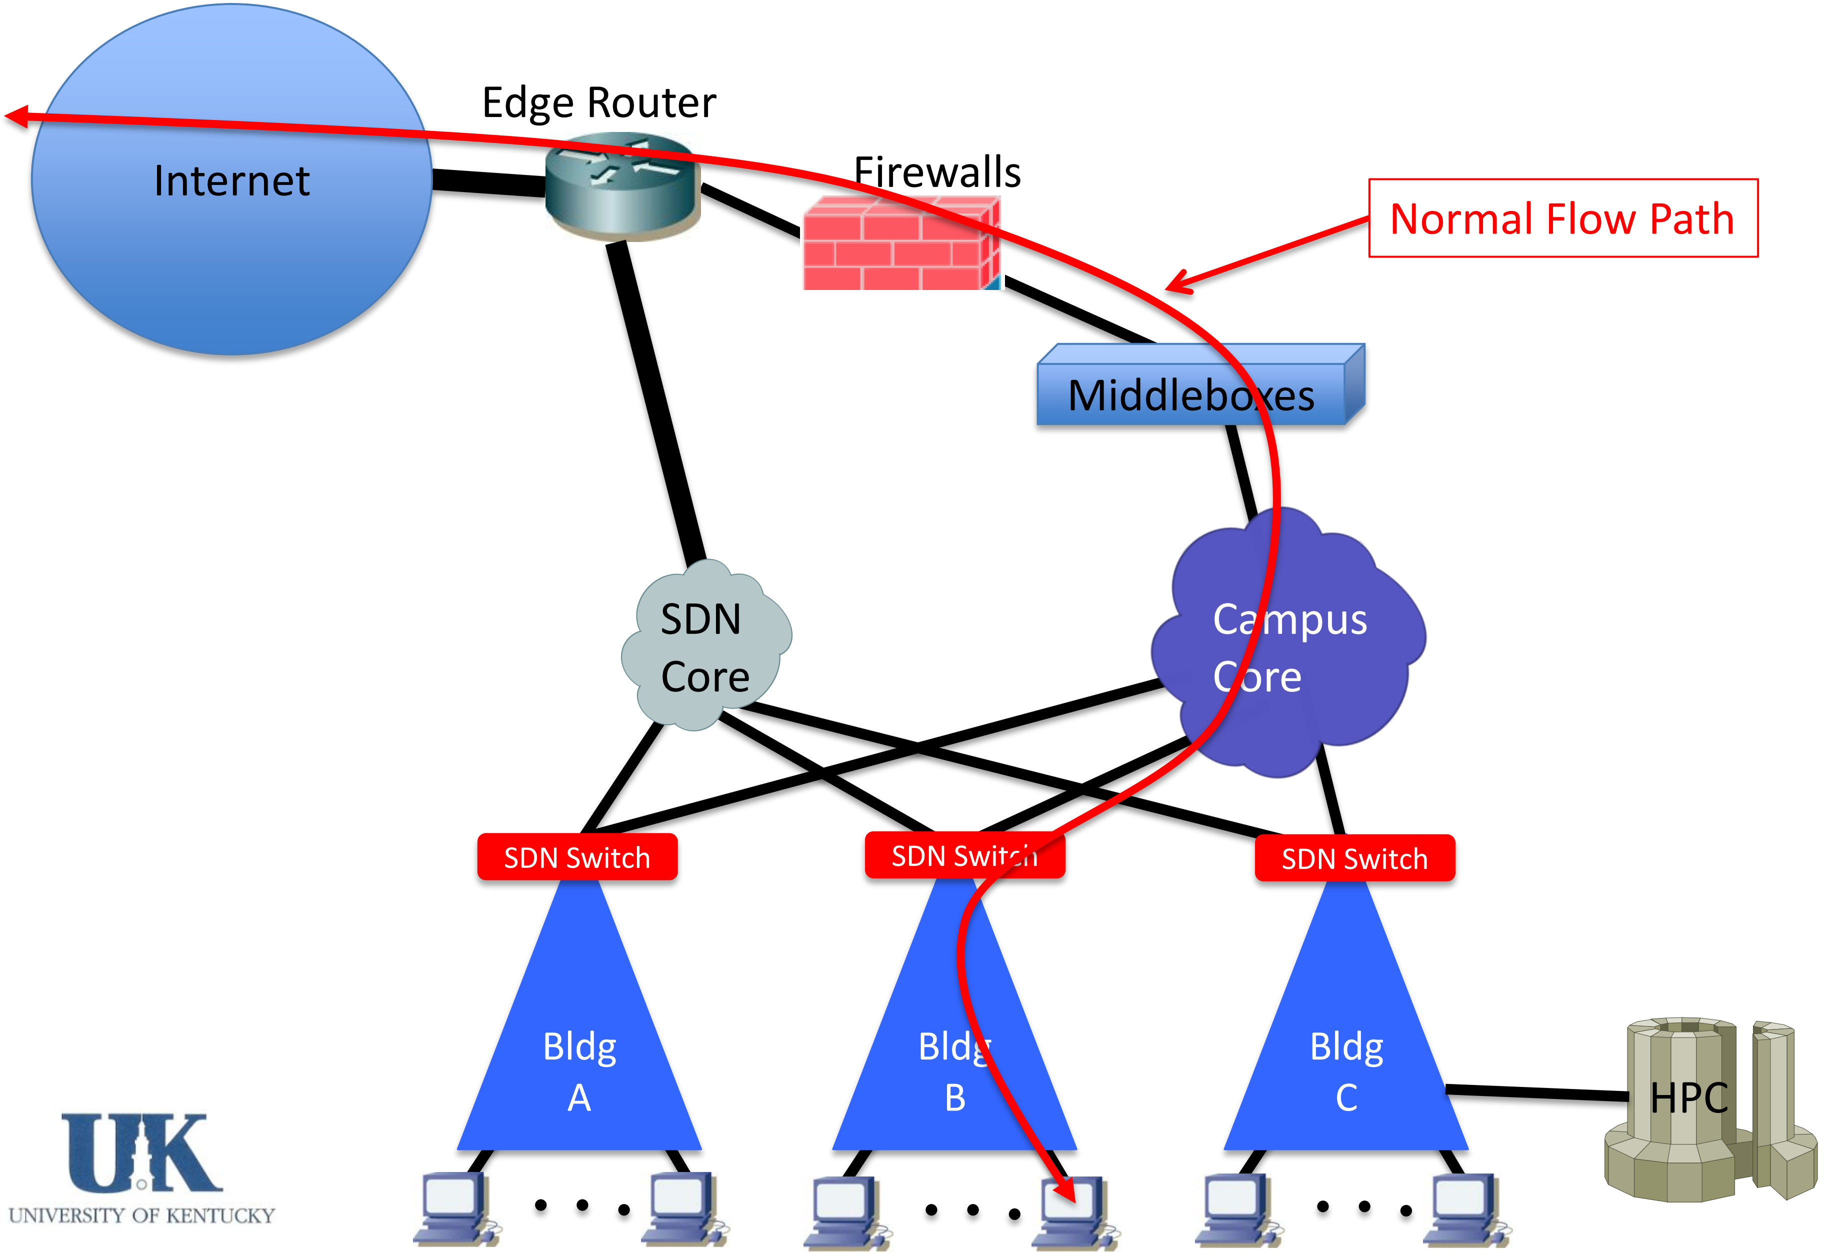

Developing an All-Campus Science DMZ

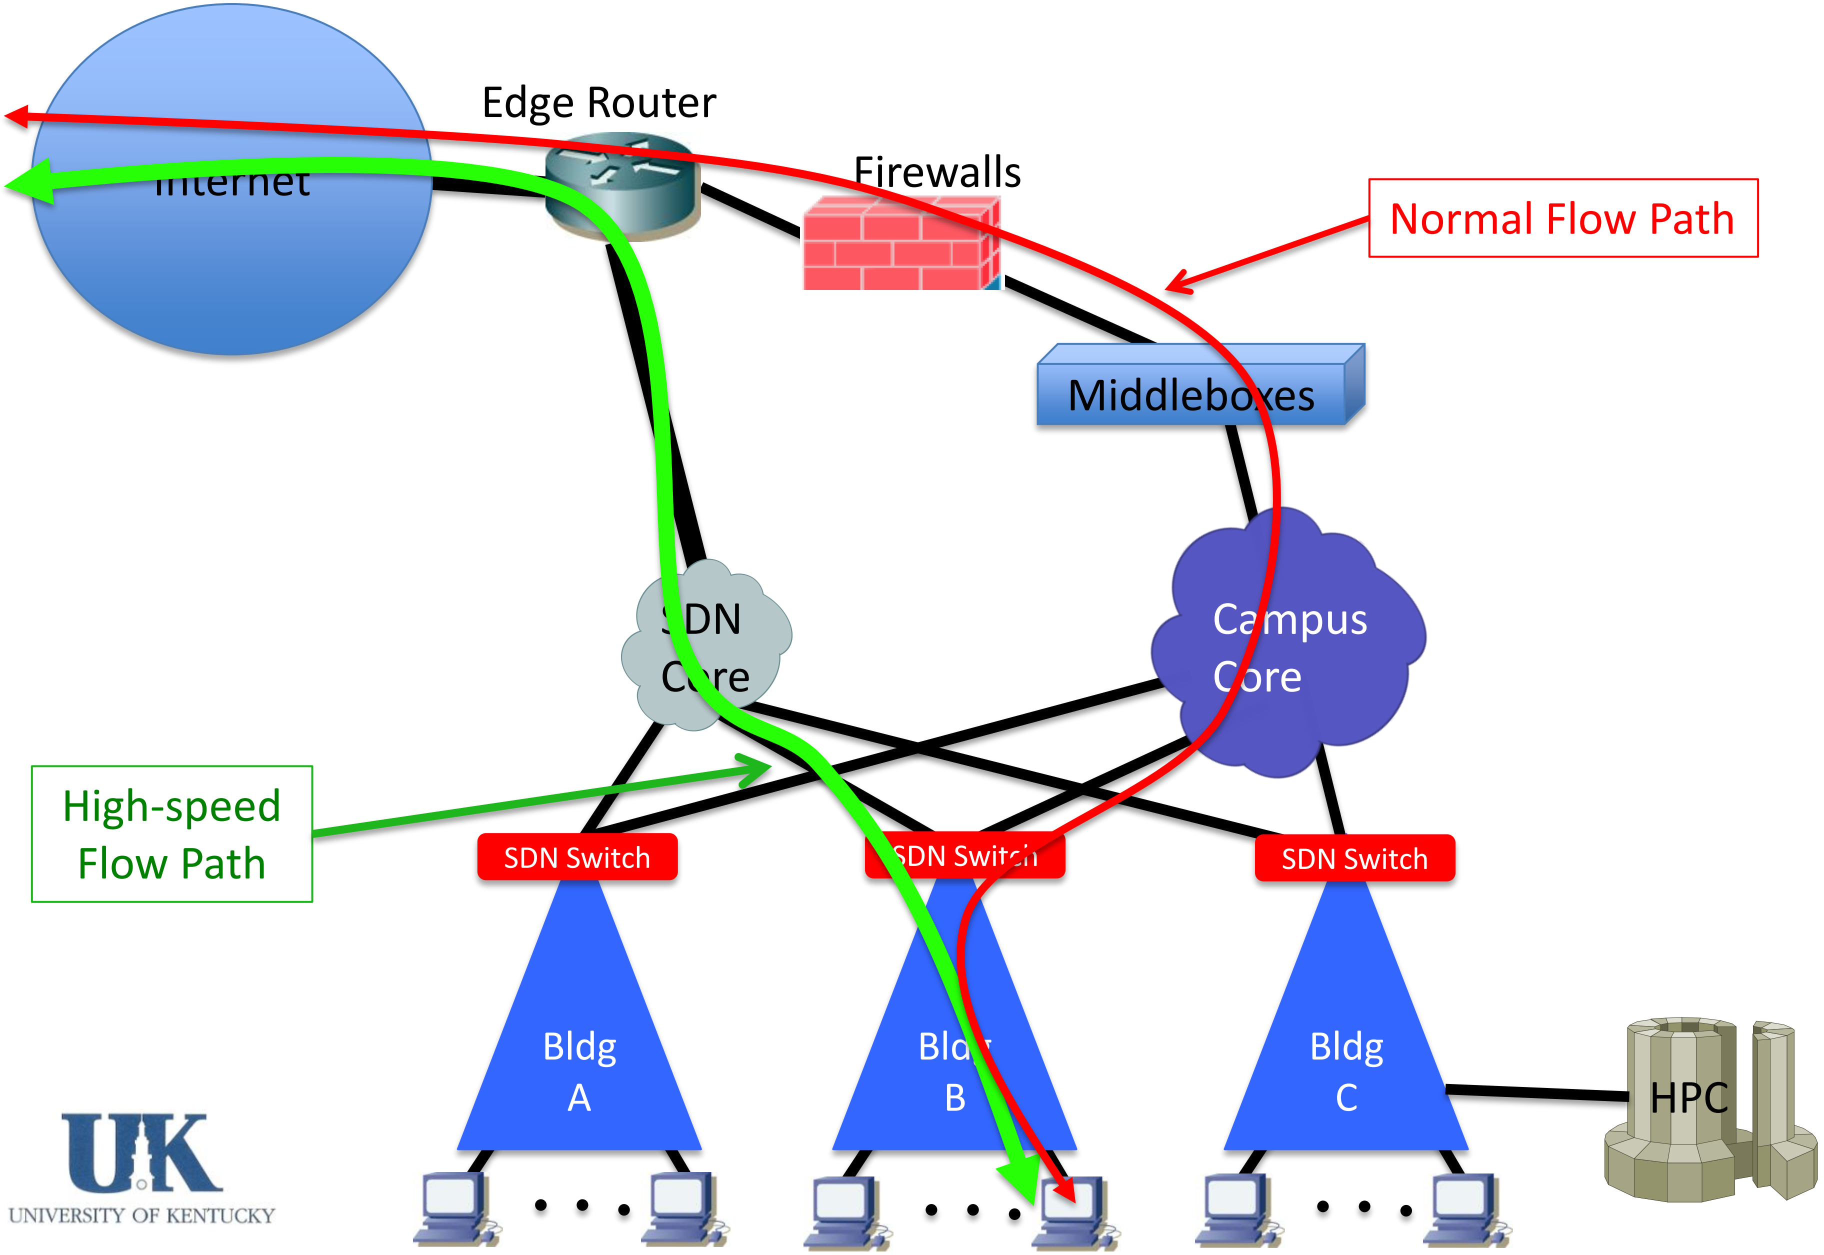

UNIVERSITY OF KENTUCKY

Developing an All-Campus Science DMZ

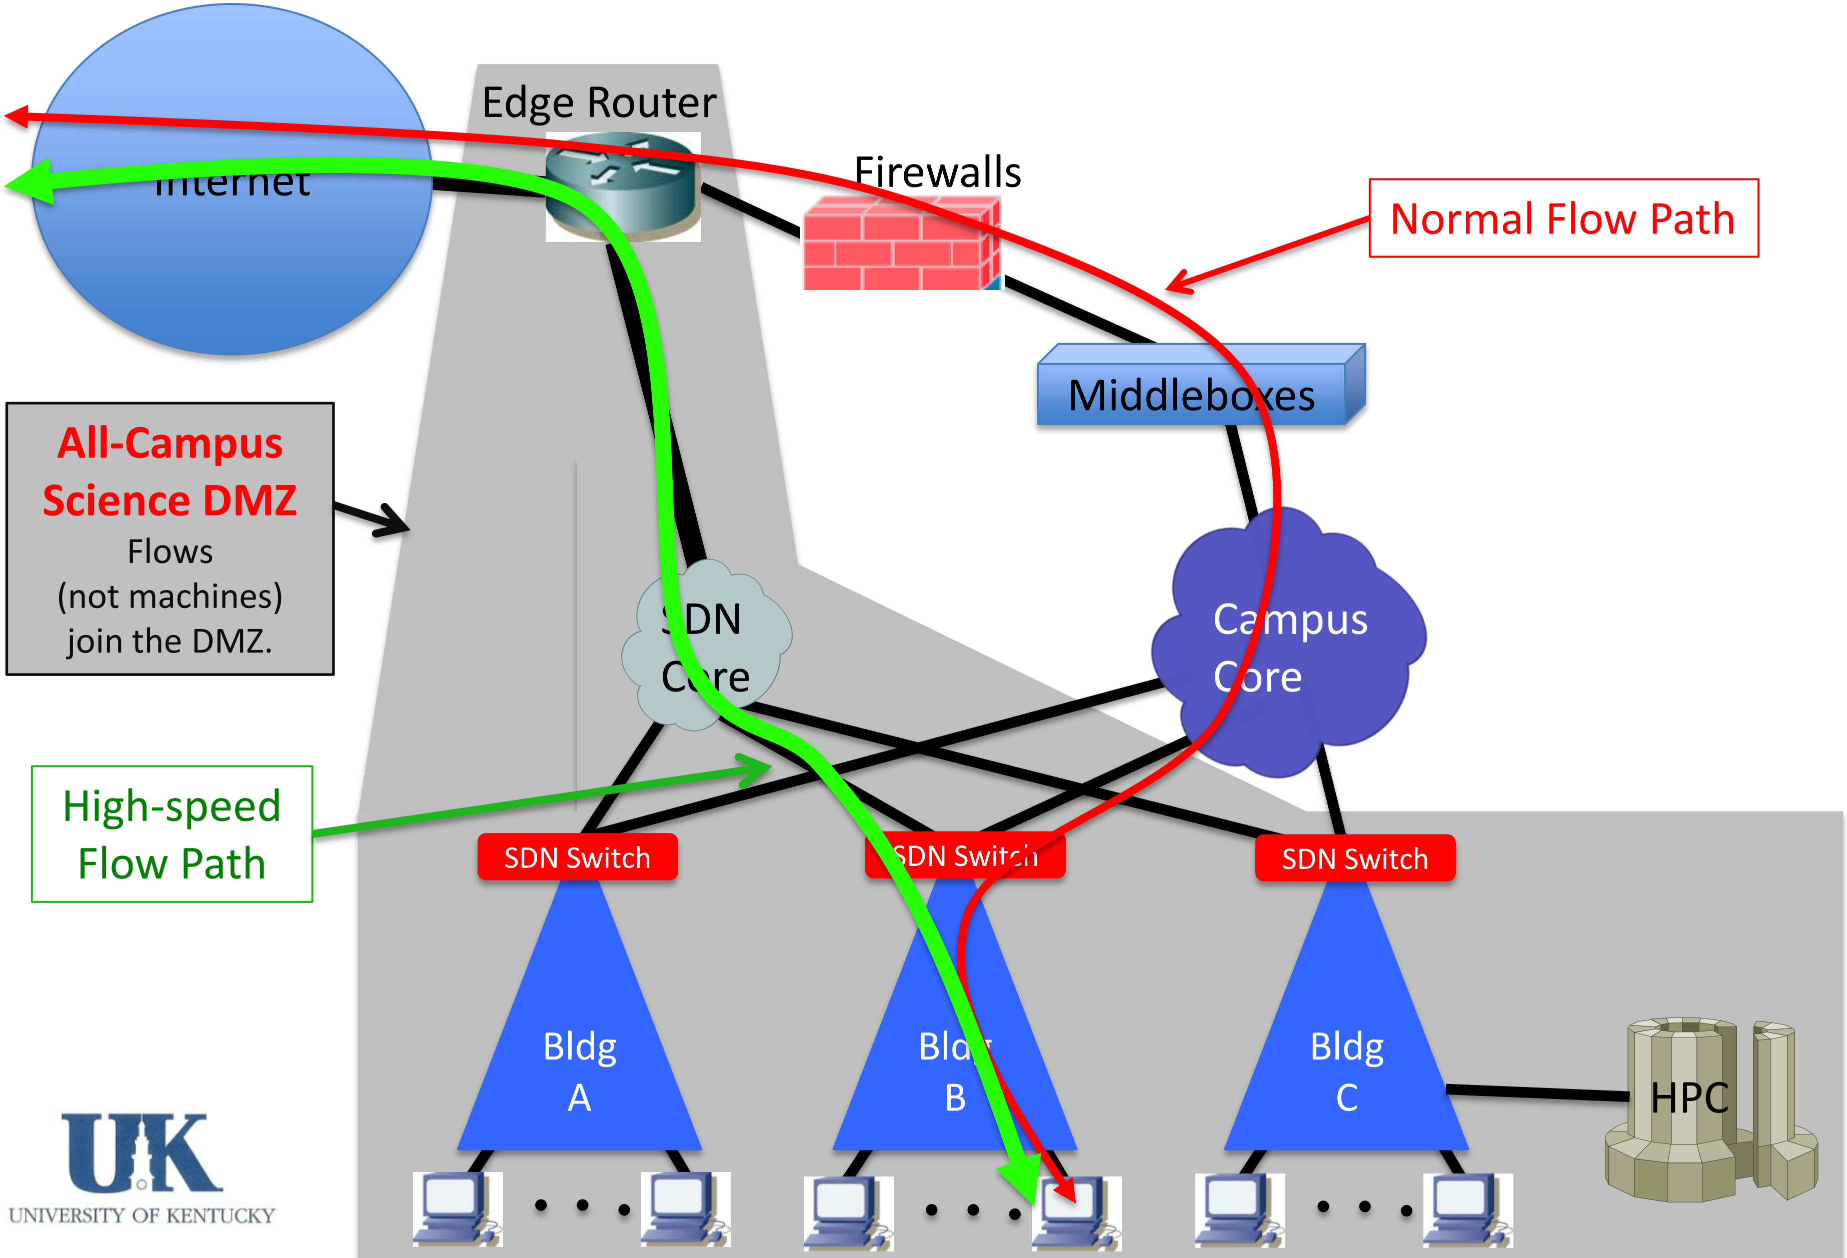

Normal Flow Path

All-Campus Science DMZ

Flows
(not machines)
join the DMZ.

High-speed Flow Path

UK
UNIVERSITY OF KENTUCKY

Caveats

- Being on the SDN network does not improve normal traffic.
- By default, traffic still routes through the slow campus network
- High-speed is only enabled for “privileged” flows
  - Must obtain permission
  - Rules must be inserted to activate the flow

Agenda

- ✓ Big data woes on the campus network
- ✓ Standard science DMZ solution
- ✓ Brief SDN overview
- ✓ A new DMZ approach
- Some results

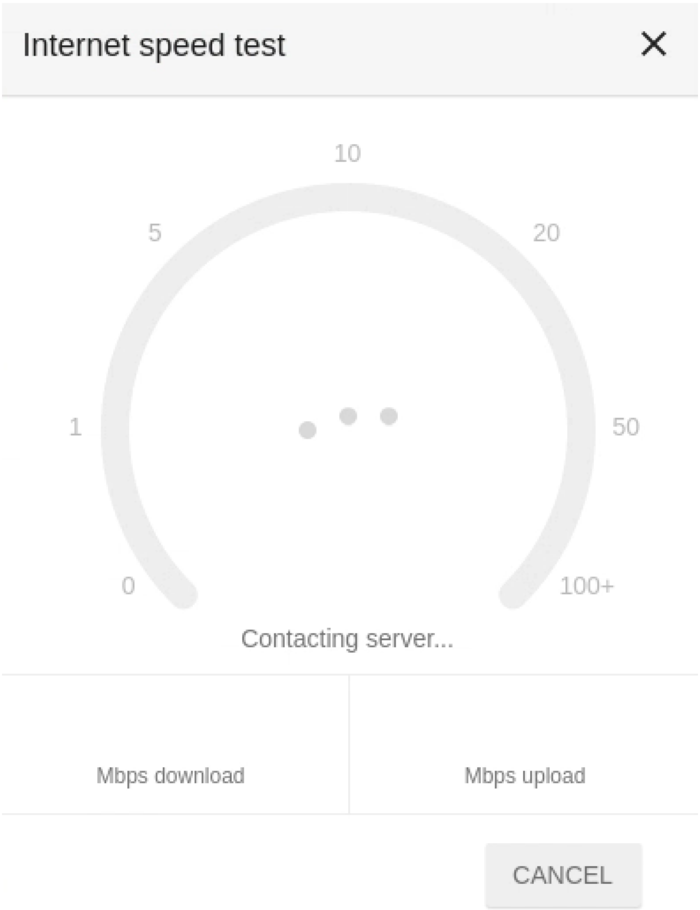

Internet speed test

Contacting server...

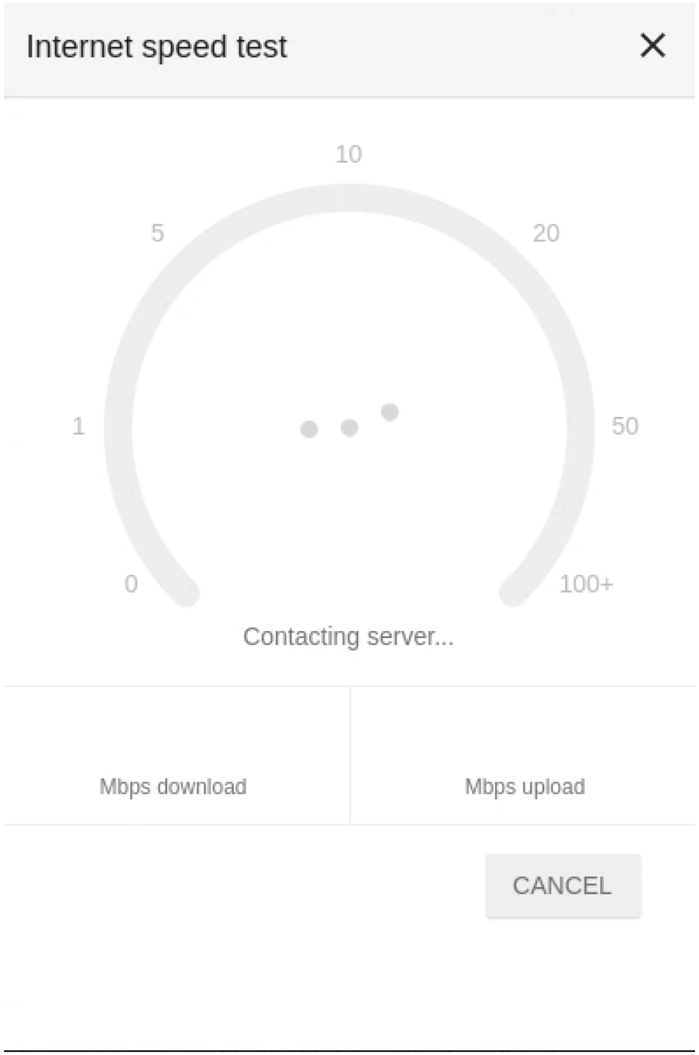

Internet speed test

Contacting server...

Perfsonar to Internet2

```
       [ ID] Interval            Transfer     Bandwidth
A) To nash-pt1.es.net
   NORMAL [ 15]   0.00-30.00  sec   168 MBytes  47.0 Mbits/sec
   SDN    [ 15]   0.00-30.00  sec  14.3 GBytes  4094 Mbits/sec

B) To atla-pt1.es.net
   NORMAL [ 15]   0.00-30.00  sec   189 MBytes  52.9 Mbits/sec
   SDN    [ 15]   0.00-30.00  sec  16.1 GBytes  4600 Mbits/sec

C) To wash-pt1.es.net
   NORMAL [ 15]   0.00-30.00  sec   282 MBytes  78.9 Mbits/sec
   SDN    [ 15]   0.00-30.00  sec  24.3 GBytes  6960 Mbits/sec

D) To fnal-pt1.es.net
   NORMAL [ 16]   0.00-30.00  sec   453 MBytes   127 Mbits/sec
   SDN    [ 16]   0.00-30.04  sec  34.5 GBytes  9879 Mbits/sec
```

Acknowledgements

Current and Past Principle Investigators

- James Griffioen
- Cody Bumgardner
- Zongming Fei
- Ken Calvert
- Vince Kellen

Funding Agencies

This work is supported, in part, by the National Science Foundation under NSF grants OCI-1246332, ACI-1541380, and ACI-1541426.

Infrastructure and Software Developers

- Matthew Moseley
- Sergio Rivera
- Mami Hayashida
- Charles Carpenter
- Yongwook Song
- Hussamuddin Nasir
- Andrew Groenewold
- Peter Oostema

Aruba VAN Support

- Shaun Wackerly

Thanks

Questions?

UK
UNIVERSITY OF KENTUCKY