

Peer-peer and Application-level Networking

CS 218 Fall 2003

Multicast Overlays

P2P applications

Napster, Gnutella, Robust Overlay Networks

Distributed Hash Tables (DHT)

Chord

CAN

Much of this material comes from UMASS
class slides

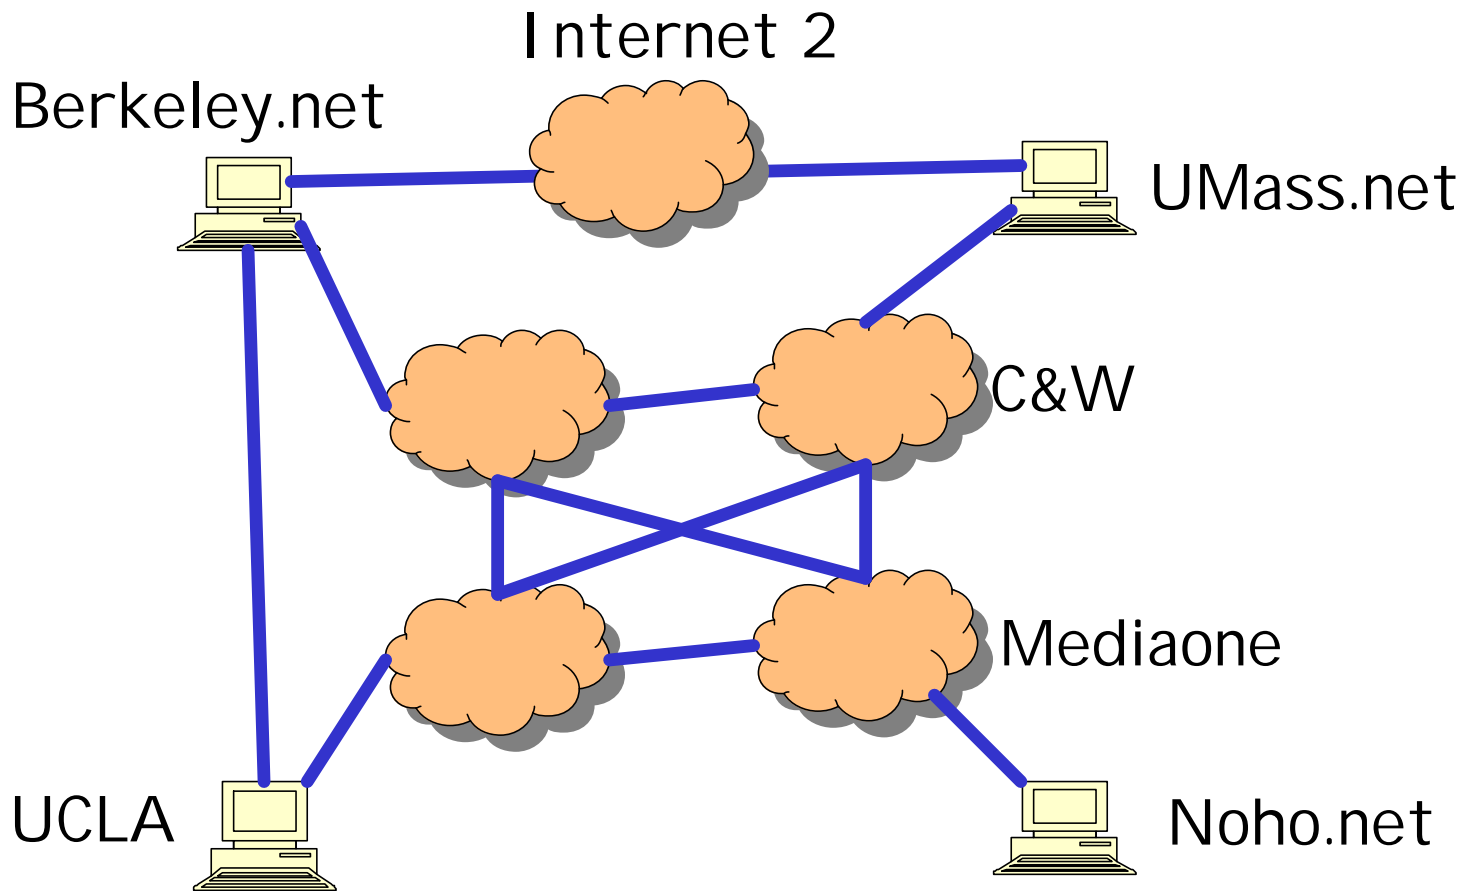
Resilient Overlay Networks

Overlay network:

- ❑ applications, running at various sites
- ❑ create “logical” links (e.g., TCP or UDP connections) pairwise between each other
- ❑ each logical link: multiple physical links, routing defined by native Internet routing
- ❑ let’s look at an example:


Internet Routing

- BGP defines routes between stub networks




Internet Routing

- BGP defines routes between stub networks




Internet Routing

- BGP defines routes between stub networks



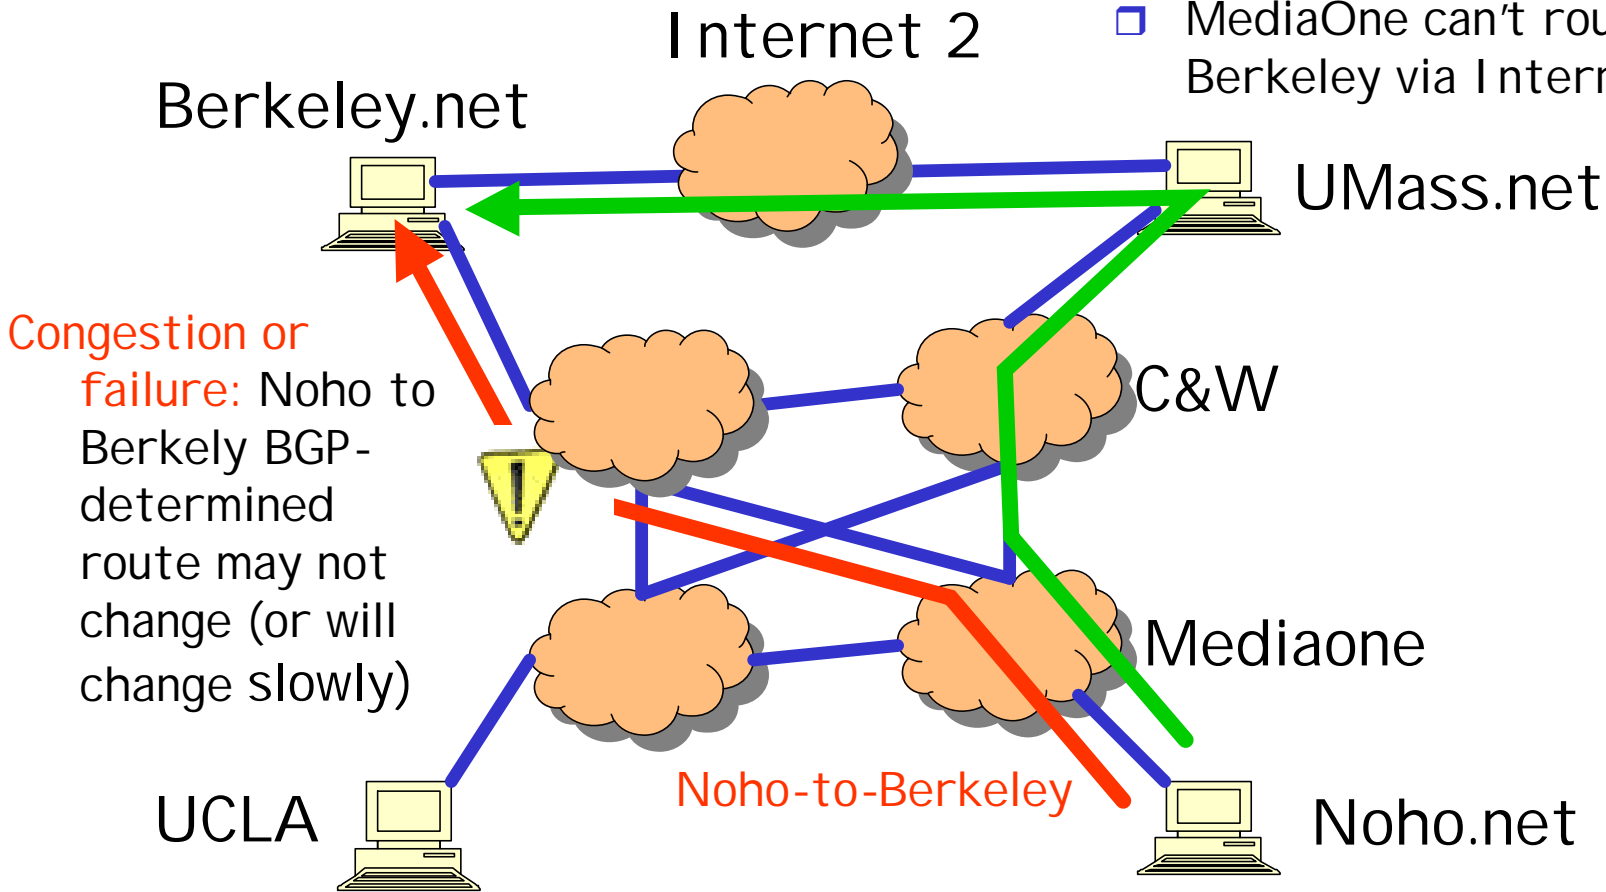
Internet Routing



Internet Routing


Noho to UMass to Berkeley

- ❑ route not visible or available via BGP!
- ❑ MediaOne can't route to Berkeley via Internet2



RON: Resilient Overlay Networks

Premise: by building application overlay network, can increase performance, reliability of routing



RON Experiments

- ❑ Measure loss, latency, and throughput with and without RON
- ❑ 13 hosts in the US and Europe
- ❑ 3 days of measurements from data collected in March 2001
- ❑ 30-minute average loss rates
 - A 30 minute outage is very serious!
- ❑ Note: Experiments done with “No-Internet2-for-commercial-use” policy

An order-of-magnitude fewer failures

30-minute average loss rates

Loss Rate	RON Better	No Change	RON Worse
10%	479	57	47
20%	127	4	15
30%	32	0	0
50%	20	0	0
80%	14	0	0
100%	10	0	0

6,825 “path hours” represented here

12 “path hours” of essentially complete outage

76 “path hours” of TCP outage

RON routed around all of these!

One indirection hop provides almost all the benefit!

RON Research Issues

- How to design overlay networks?
 - Measurement and self-configuration
 - Understanding performance of underlying net.
 - Fast fail-over.
 - Sophisticated metrics.
 - application-sensitive (e.g., delay versus throughput) path selection.
- Effect of RON on underlying network
 - If everyone does RON, are we better off?