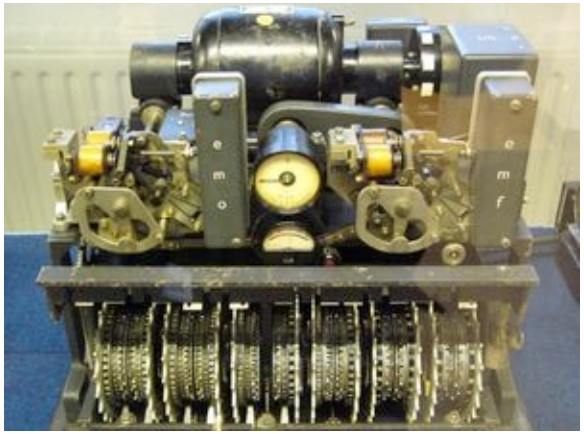

Class Meeting: Lectures 21-23

COS 461: Computer Networks

Kyle Jamieson

Today

- Network Security
- Datacenter Networks
- Course Summary & Wrap-Up

BGP Security Today

- Applying "best common practices"
  - Securing the session (authentication, encryption)
  - Filtering routes by prefix and AS path
  - Packet filters to block unexpected control traffic

- This is not good enough
  - Depends on vigilant application of practices
  - Doesn't address fundamental problems
    - Can't tell who owns the IP address block
    - Can't tell if the AS path is bogus or invalid
    - Can't be sure data packets follow the chosen route

Proposed Enhancements to BGP

Secure BGP

Origin Authentication + cryptographic signatures

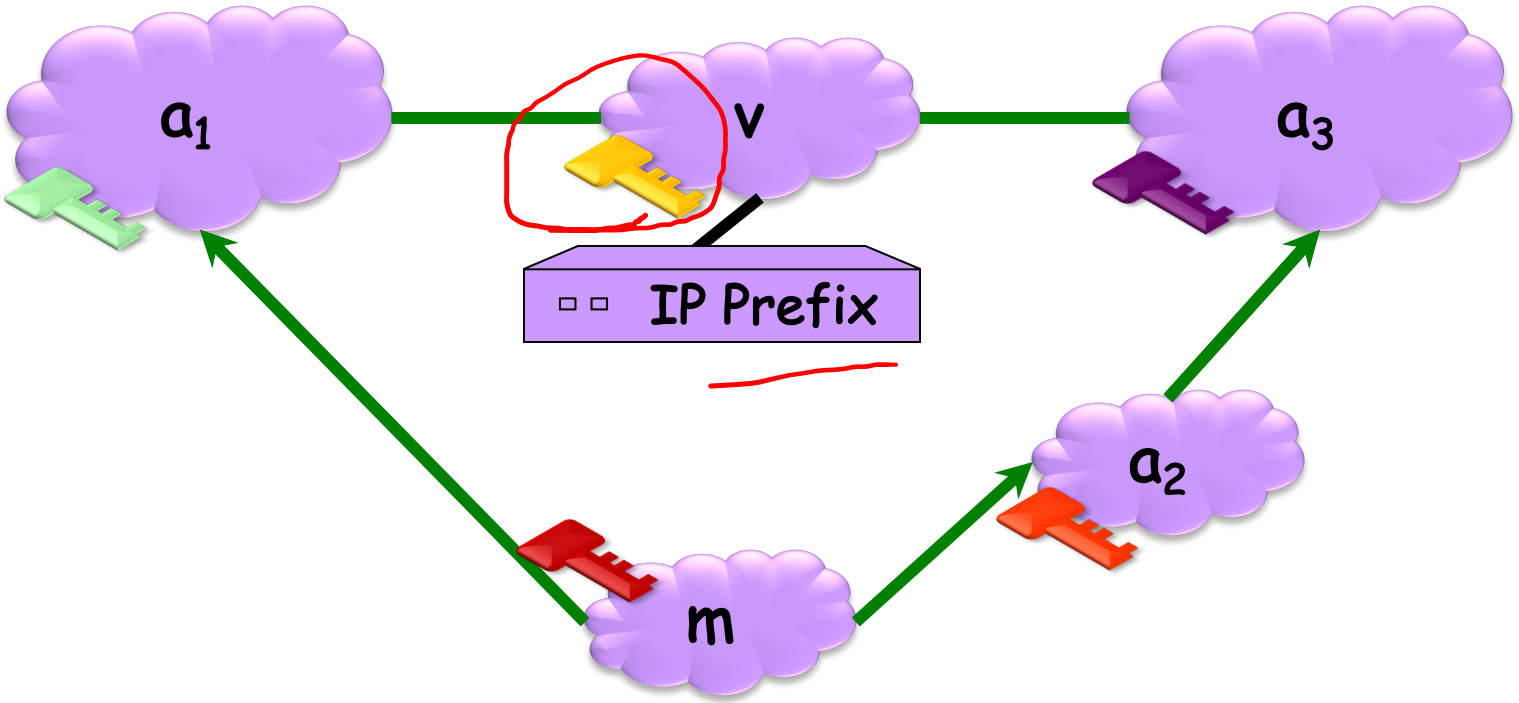

Public Key Signature: Anyone who knows v's public key can verify that the message was sent by v.

Secure BGP

Origin Authentication + cryptographic signatures

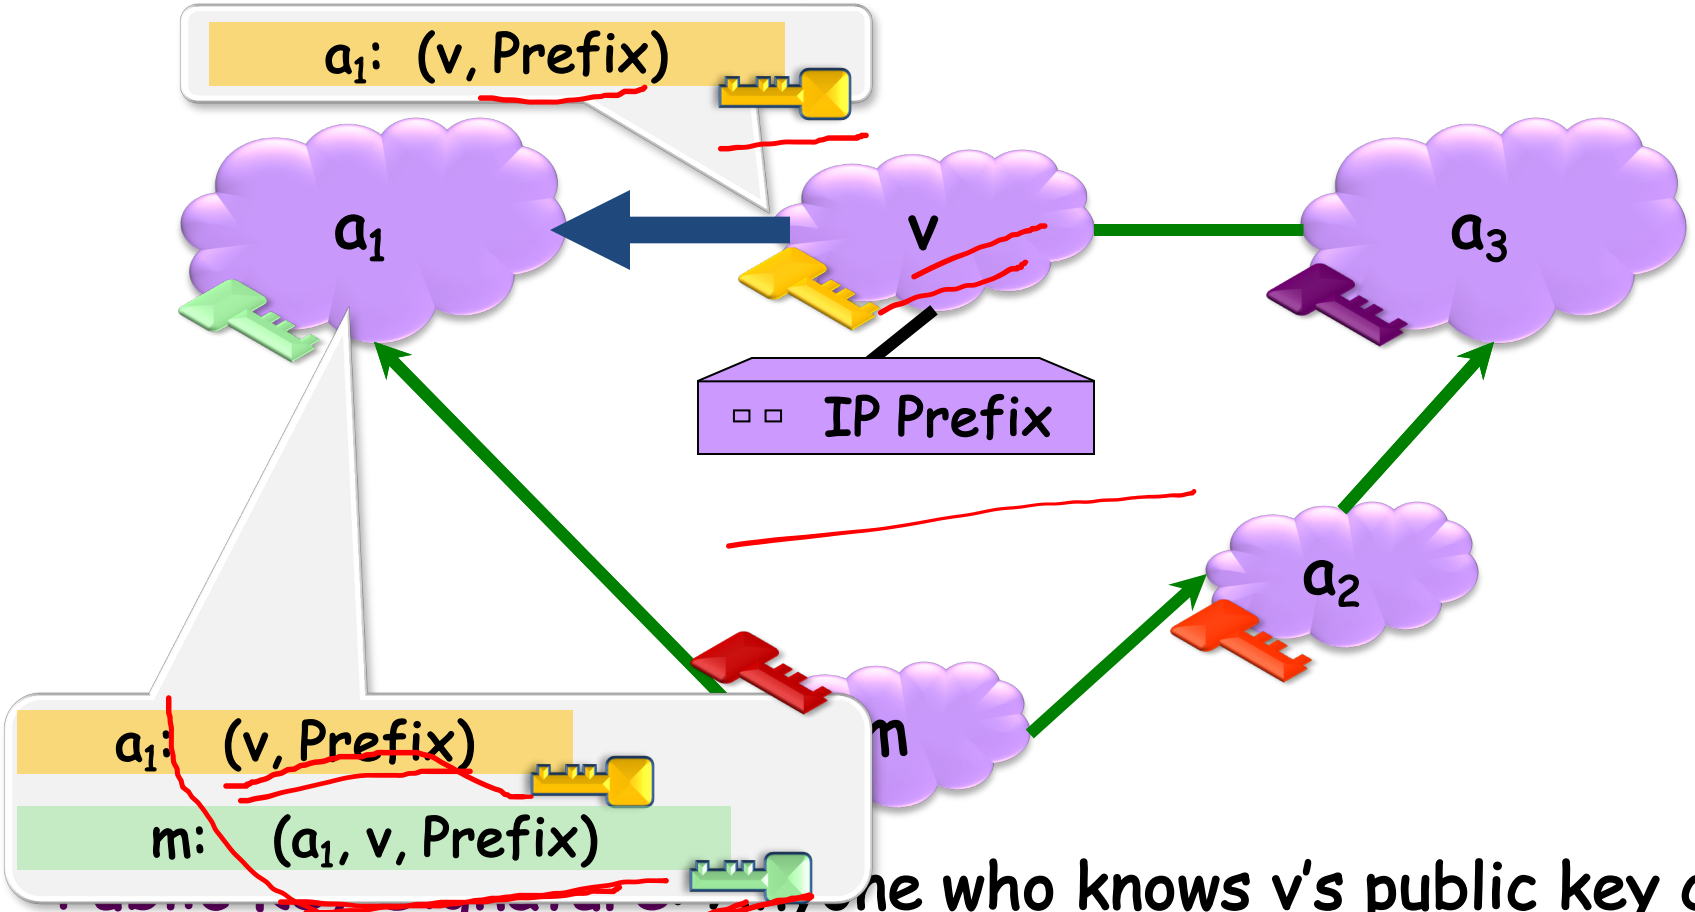

Anyone who knows v's public key can verify that the message was sent by v.

"Secure BGP"

- Route attestations
  - Distributed as an attribute in BGP update message
  - Signed by each AS as route traverses the network
- Address attestations
  - Claim the right to originate a prefix
  - Signed and distributed out-of-band
  - Checked through delegation chain from ICANN
- S-BGP can validate
  - AS path indicates the order ASes were traversed
  - No intermediate ASes were added or removed
  - Proper ASes originate prefixes

S-BGP Deployment Challenges

- Complete, accurate registries of prefix "owner"
- Public Key Infrastructure
  - To know the public key for any given AS
- Cryptographic operations
  - E.g., digital signatures on BGP messages
- Need to perform operations quickly
  - To avoid delaying response to routing changes
- Difficulty of incremental deployment
  - Hard to have a "flag day" to deploy S-BGP

Detecting Suspicious Routes

- Monitoring BGP update messages
  - Use past history as an implicit registry
- E.g., AS that announces each address block
  - Prefix 18.0.0.0/8 usually originated by AS 3
- E.g., AS-level edges and paths
  - Never seen the subpath "7018 88 1785"
- Out-of-band detection mechanism
  - Generate reports and alerts
  - Internet Alert Registry: http://iar.cs.unm.edu/
  - Prefix Hijack Alert System: http://phas.netsec.colostate.edu/

BGP Security: Conclusions

- Internet protocols designed based on trust
  - Insiders are good actors, bad actors on the outside
- Border Gateway Protocol is very vulnerable
  - Glue that holds the Internet together
  - Hard for an AS to locally identify bogus routes
  - Attacks can have serious global consequences
- Proposed solutions/approaches
  - Secure variants of the Border Gateway Protocol
  - Anomaly detection, with automated response
  - Broader focus on data-plane availability

Today

- Network Security
- **Datacenter Networks**
- Course Summary & Wrap-Up

Top-of-Rack Architecture

- Rack of servers
  - Commodity servers
  - And top-of-rack switch
- Modular design
  - Preconfigured racks
  - Power, network, and storage cabling

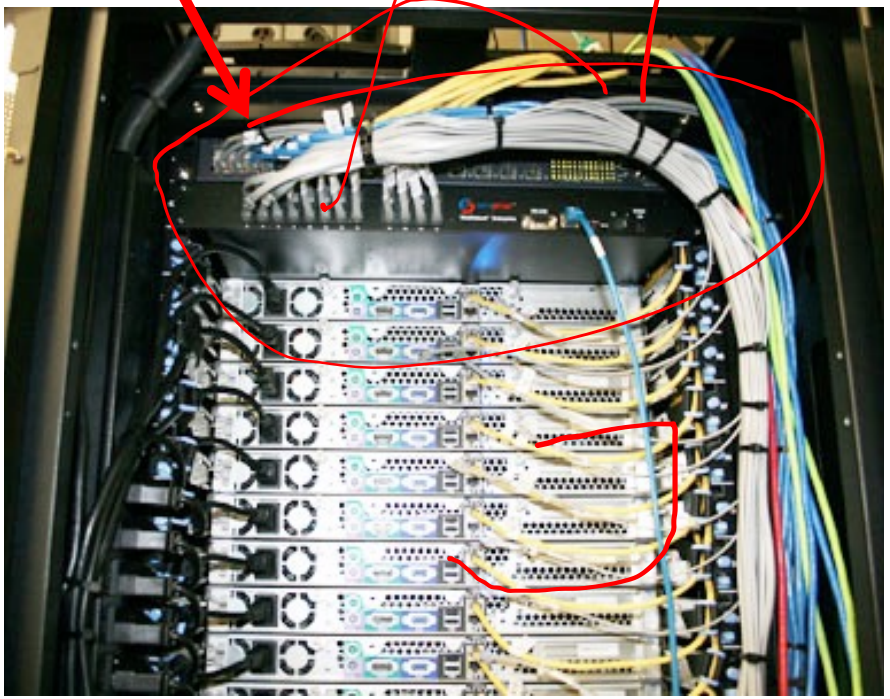

Datacenter Network Topology

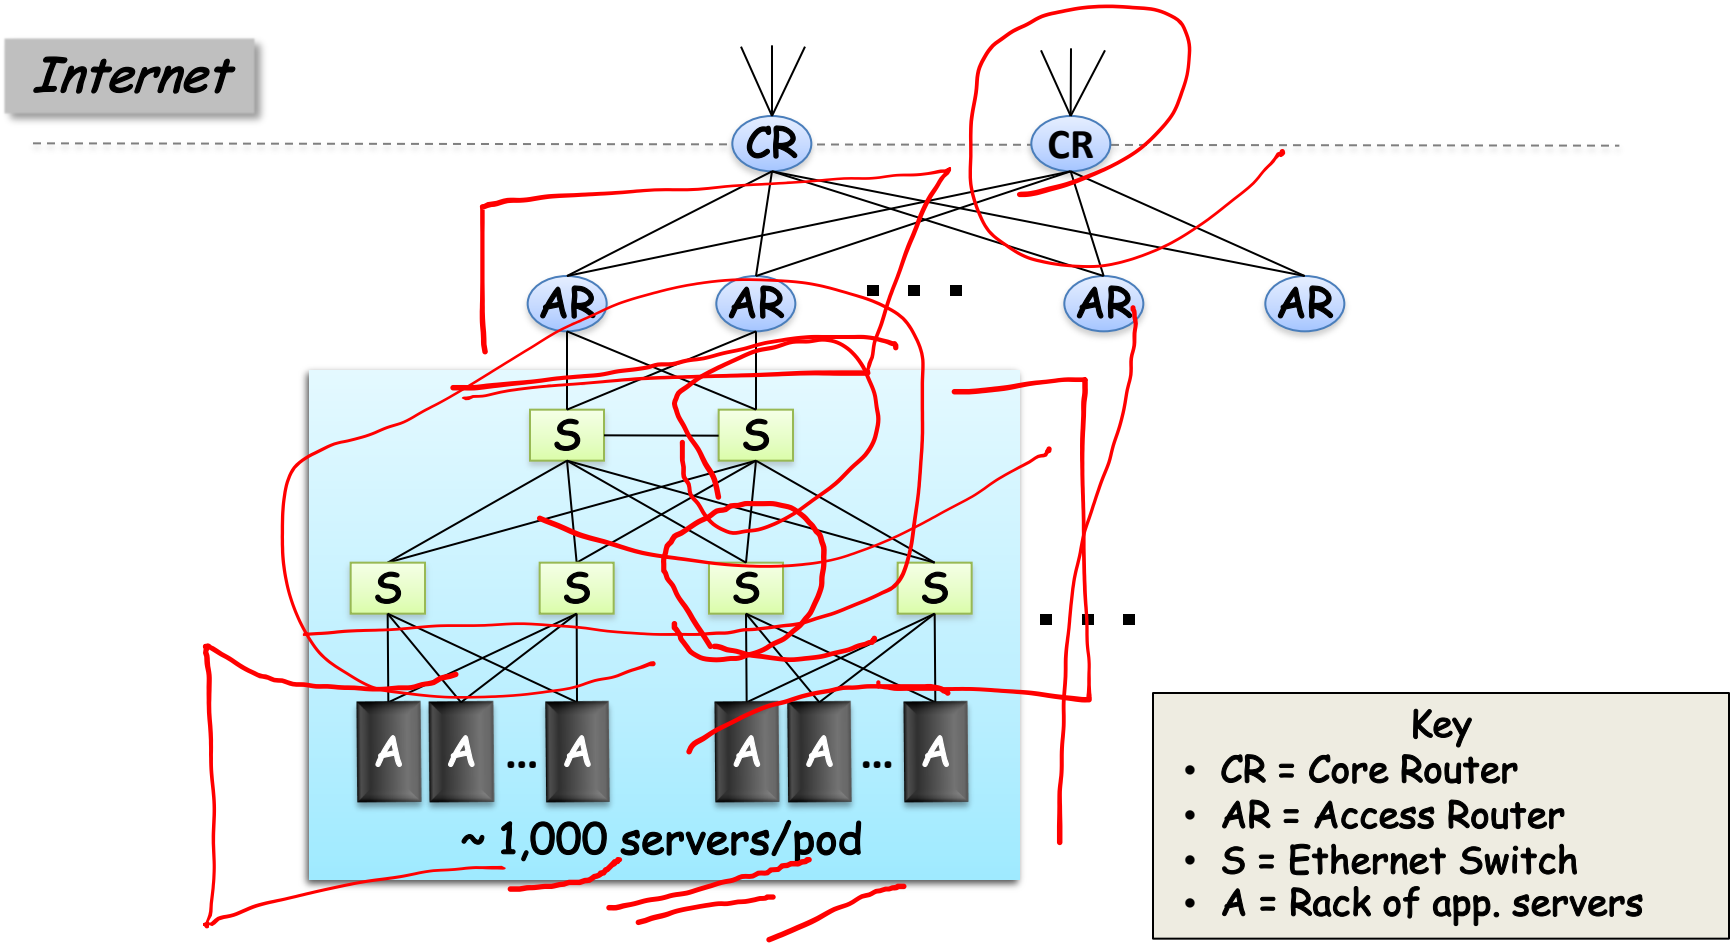

Capacity Mismatch?

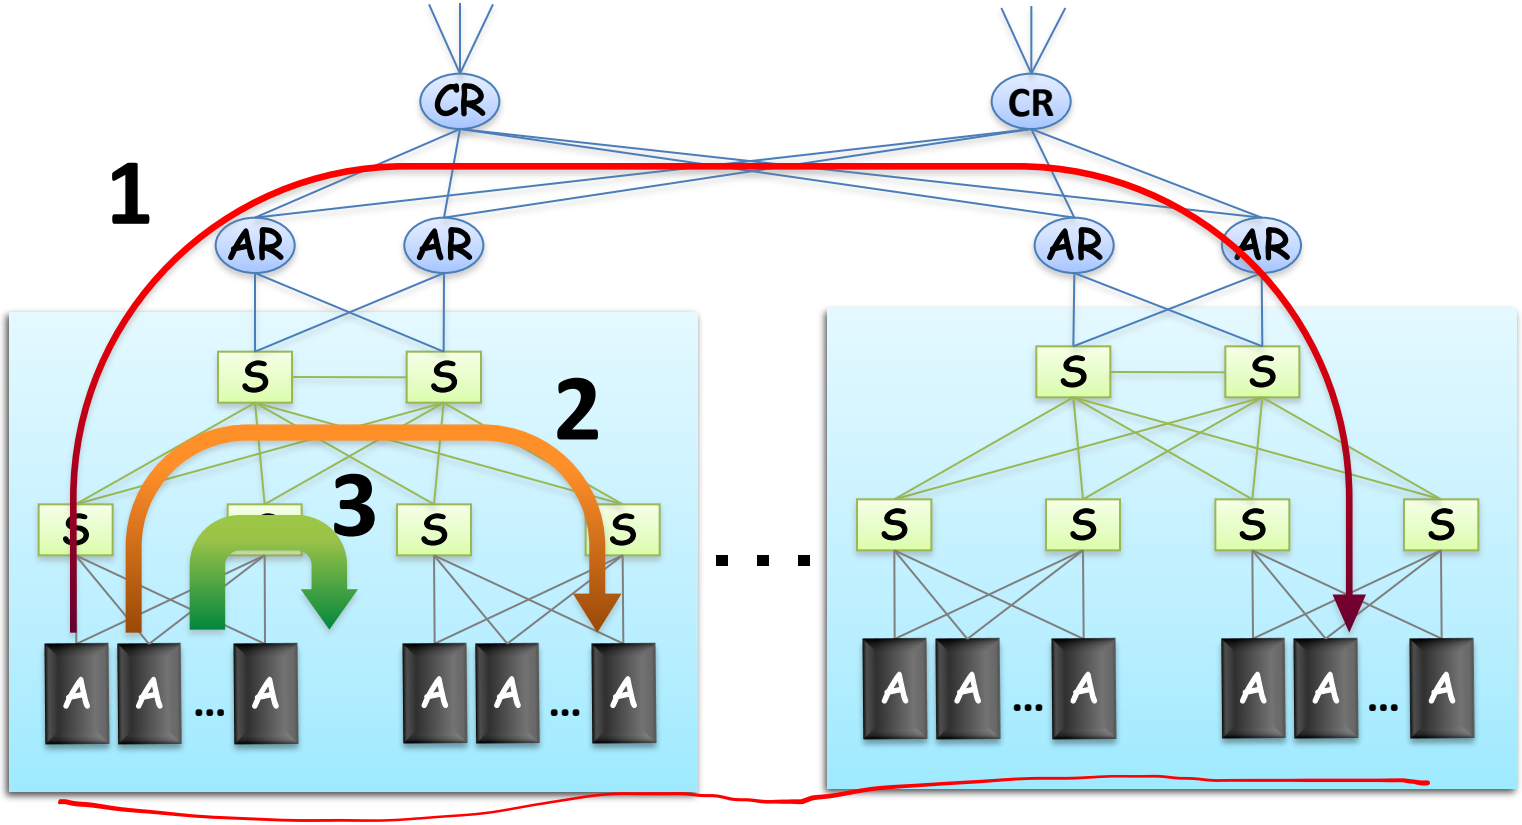

“Oversubscription”:
Much more demand vs. supply for higher links

Capacity Mismatch!

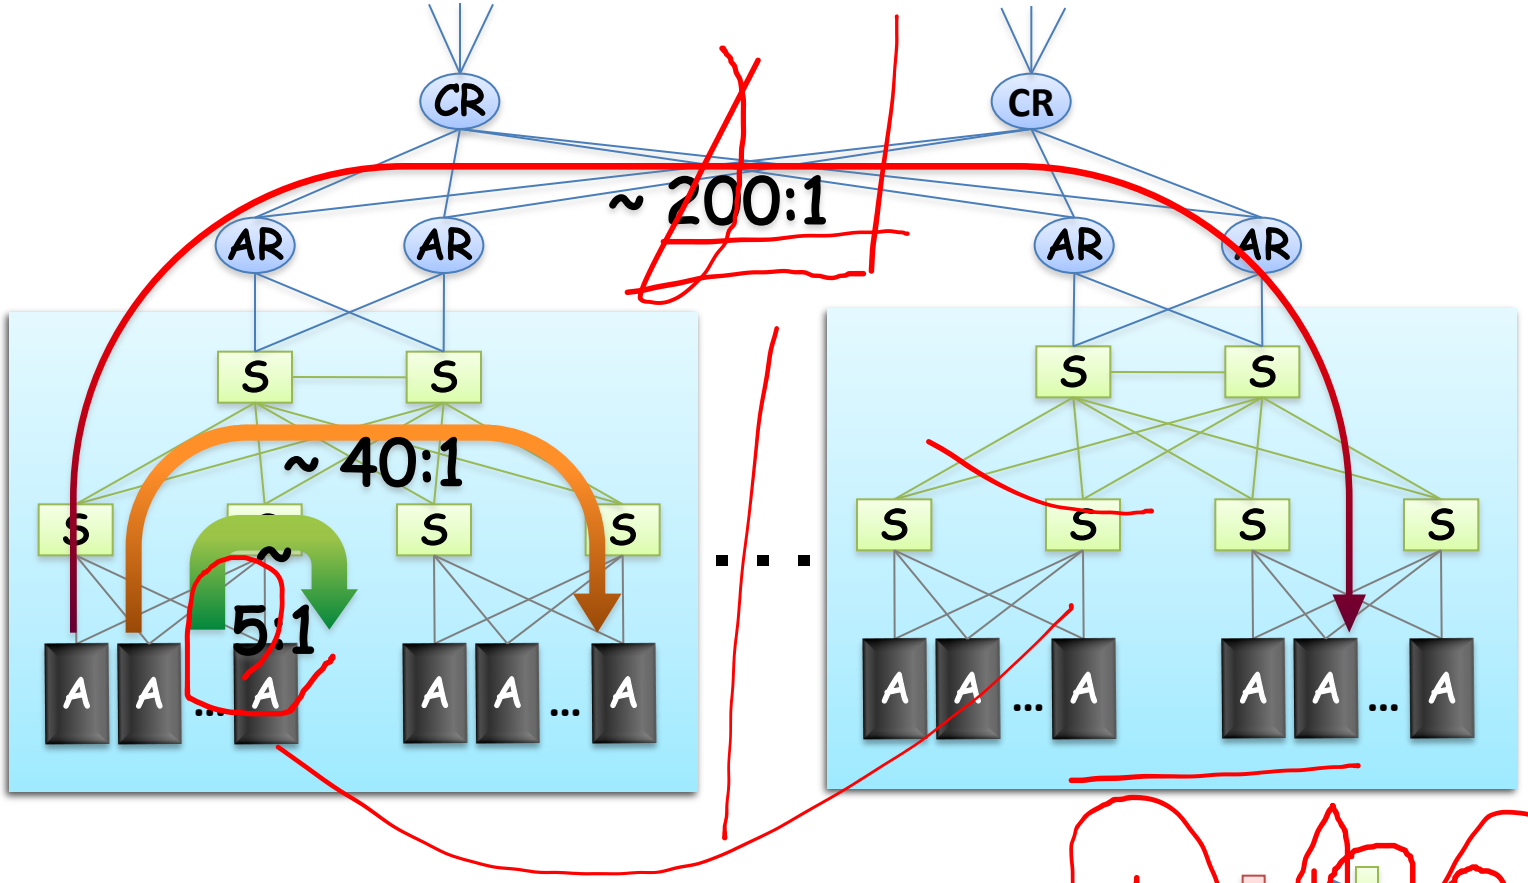

Particularly bad for east-west traffic

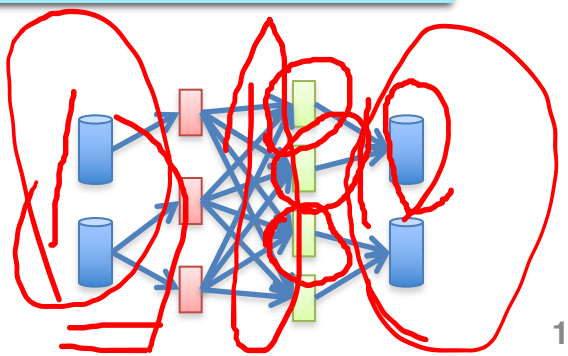

Layer 2 vs. Layer 3?

- Ethernet switching (layer 2)
  - Cheaper switch equipment
  - Fixed addresses and auto-configuration
  - Seamless mobility, migration, and failover
- IP routing (layer 3)
  - Scalability through hierarchical addressing
  - Efficiency through shortest-path routing
  - Multipath routing through equal-cost multipath

Datacenter Routing

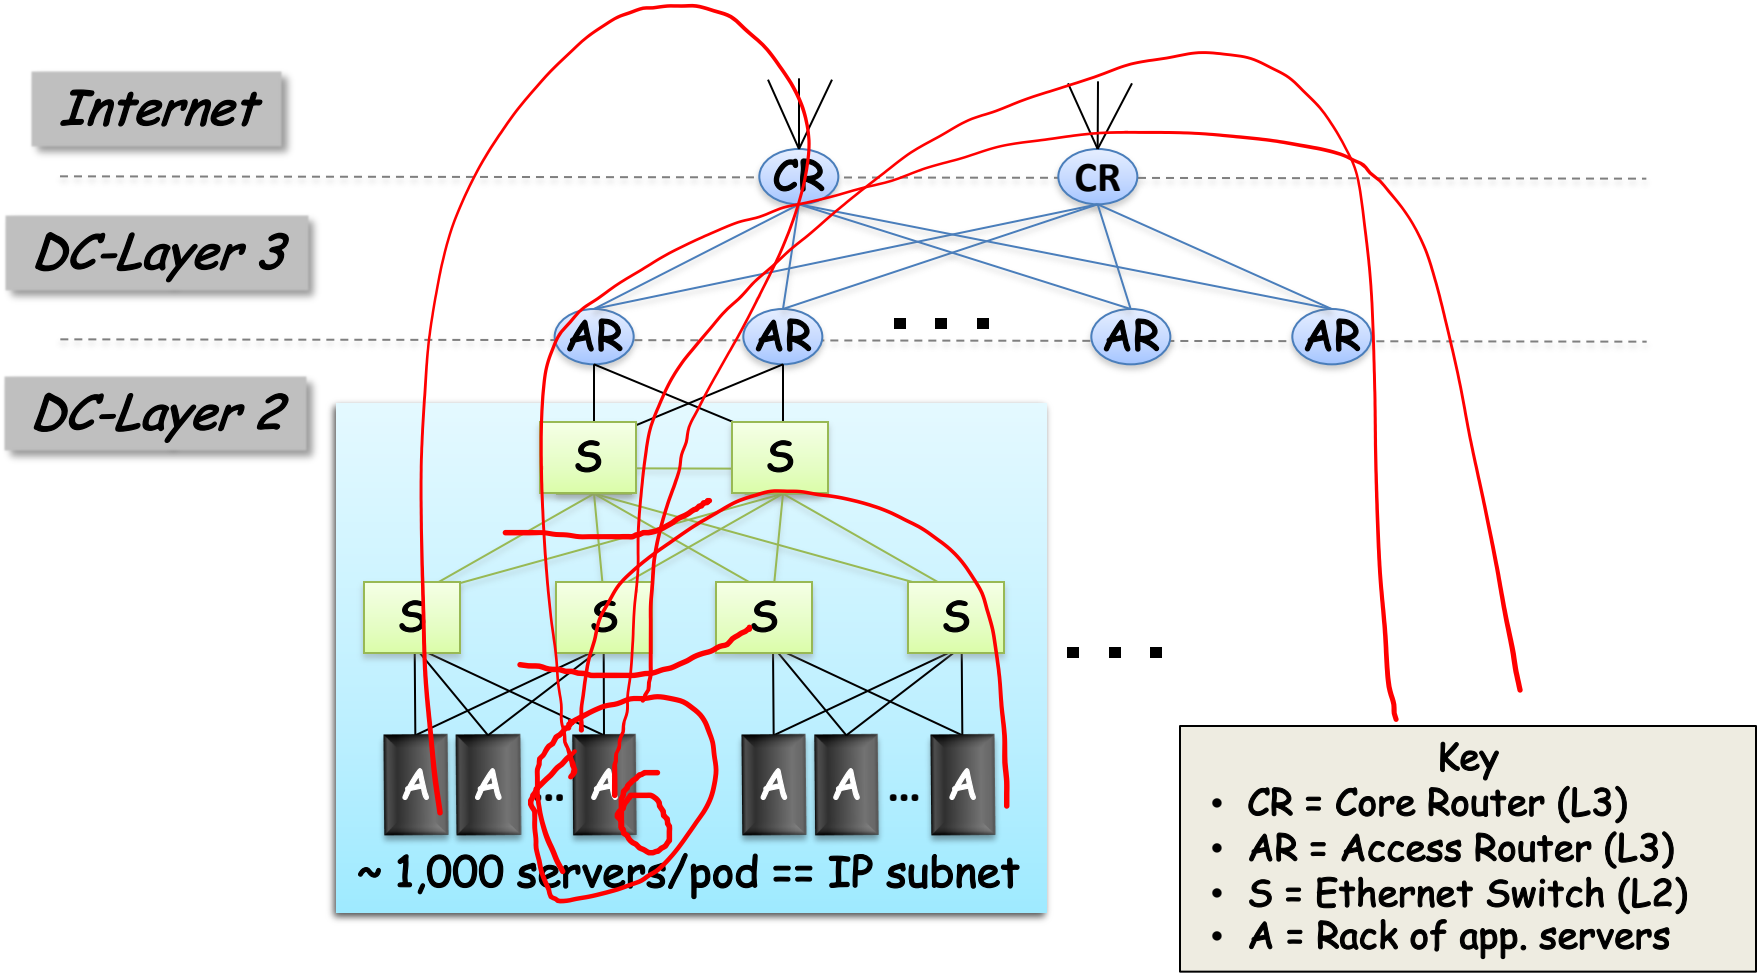

New datacenter networking problems have emerged...

Network *Incast* Problem

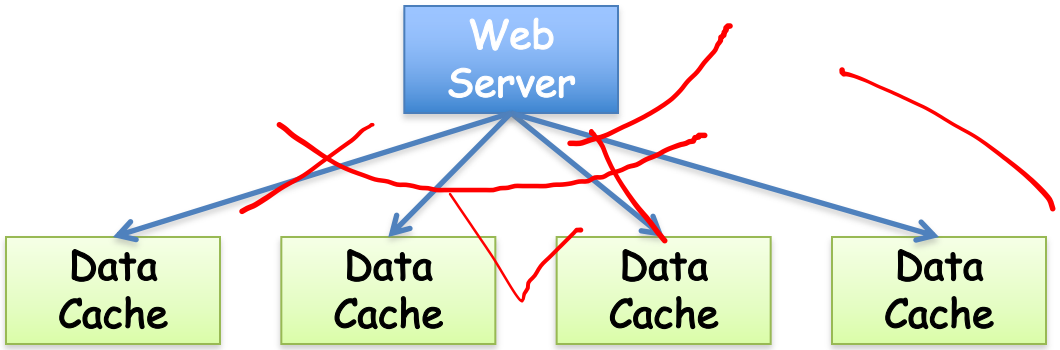

- Incast arises from synchronized, parallel requests
  - Web server sends out parallel request ("which friends of Johnny are online?"
  - Nodes reply at same time, cause traffic burst
  - Replies potentially exceed switch's buffer, causing drops

Data Center Networks: Summary

- Cloud computing
  - Major trend in IT industry
  - Today's equivalent of factories
- Datacenter networking
  - Regular topologies interconnecting VMs
  - Mix of Ethernet and IP networking
- Modular, multi-tier applications
  - New ways of building applications
  - New performance challenges

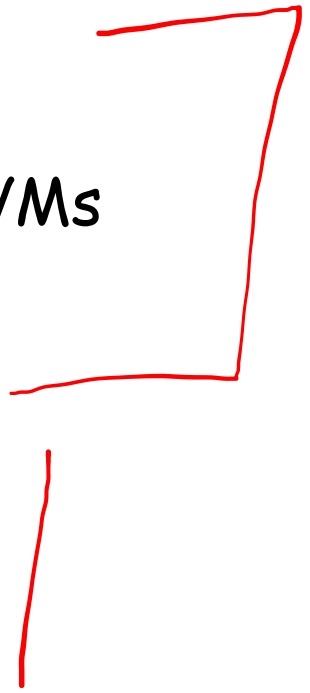

Today

- Network Security
- Datacenter Networks
- **Course Summary & Wrap-Up**

Some Key Concepts

- Course was organized around protocols
  - But a small set of concepts recur in many protocols
- General CS concepts
  - Hierarchy, indirection, caching, randomization
- Networking-specific concepts
  - Soft state, layering, (de)multiplexing
  - End-to-end argument

Hierarchy

- Scalability of large systems
  - Cannot store all information everywhere
  - Cannot centrally coordinate everything
- Hierarchy to manage scale
  - Divide system into smaller pieces
- Hierarchy to divide control
  - Decentralized management
- Examples from the Internet
  - IP addresses, routing protocols, DNS, P2P

Indirection

- Referencing by name
  - Rather than the value itself
  - E.g., manipulating a variable through a pointer
- Benefits of indirection
  - Human convenience
  - Reducing overhead when things change
- Examples of indirection in the Internet
  - Names vs. addresses
  - Mobile IP

Caching

- Duplicating data stored elsewhere
  - To reduce latency for accessing the data
  - To reduce resources consumed
- Caching is often quite effective
  - Speed difference between cache and primary copy
  - Locality of reference, and small set of popular data
- Examples from the Internet
  - DNS caching, Web caching, CDNs

Randomization

- Distributed adaptive algorithms
  - Multiple distributed parties
  - Adapting independently
- Risk of synchronization
  - Many parties reacting at the same time
  - Leading to bad aggregate behavior
- Randomization can desynchronize
  - Ethernet back-off
- Rather than imposing centralized control

Soft State

- State: stored in nodes by network protocols
  - Installed by receiver of a set-up message
  - Updated when conditions change

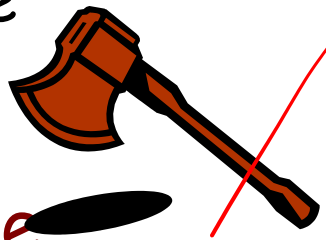

- Hard state: valid unless told otherwise
  - Removed by receiver of tear-down message
  - Requires error handling to deal with sender failure
- Soft state: invalid if not told to refresh
  - Periodically refreshed, removed by timeout

- Soft state reduces complexity
  - DNS caching, DHCP leases

Layering: A Modular Approach

- Sub-divide the problem
  - Each layer relies on services from layer below
  - Each layer exports services to layer above
- Interface between layers defines interaction
  - Hides implementation details
  - Layers can change without disturbing other layers

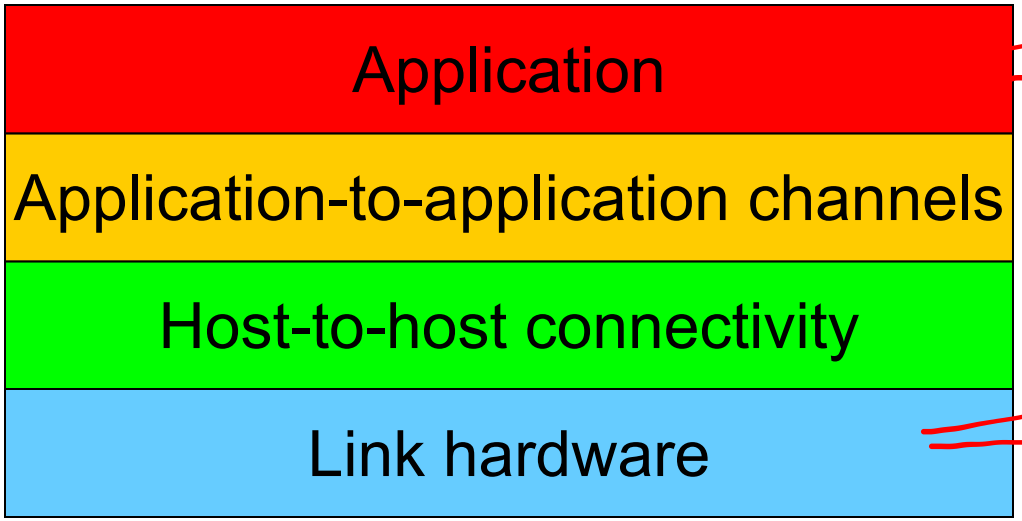

Power at the End Host

End-to-End Principle

Whenever possible, communications protocol operations should be defined to occur at the end-points of a communications system.

Programmability

With programmable end hosts, new network services can be added at any time, by anyone.

The Internet of the Future

- Can we fix what ails the Internet
  - Security, performance, reliability
  - Upgradability, managability
- Without throwing out baby with bathwater
  - Ease of adding new hosts
  - Ease of adding new services
  - Ease of adding new link technologies
- An open technical and policy question…

Final Exam

- **Begins 9:00 AM on Wednesday, December 15.**
- The exam will be due at **5:00 PM on Monday, December 20**
  - **Six hour** Gradescope completion time limit
- online, open-book, open-461 material, calculators-allowed