MAC6958
Tópicos Avançados em Ciência de Dados para Redes de Computadores

- **Prof. Daniel Macêdo Batista**
  Prof. Roberto Hirata
- **DCC-IME-USP**

  - 24/3/2020

Roteiro

- **BGP** (Alan Barzilay)
- **Processamento de logs** (Giovana Vieira de Morais)

Internet inter-AS routing: BGP

- BGP (Border Gateway Protocol): *the* de facto inter-domain routing protocol
  - "glue that holds the Internet together"
- BGP provides each AS a means to:
  - eBGP: obtain subnet reachability information from neighboring ASes
  - iBGP: propagate reachability information to all AS-internal routers.
  - determine "good" routes to other networks based on reachability information and *policy*
- allows subnet to advertise its existence to rest of Internet: *"I am here"*

eBGP, iBGP connections

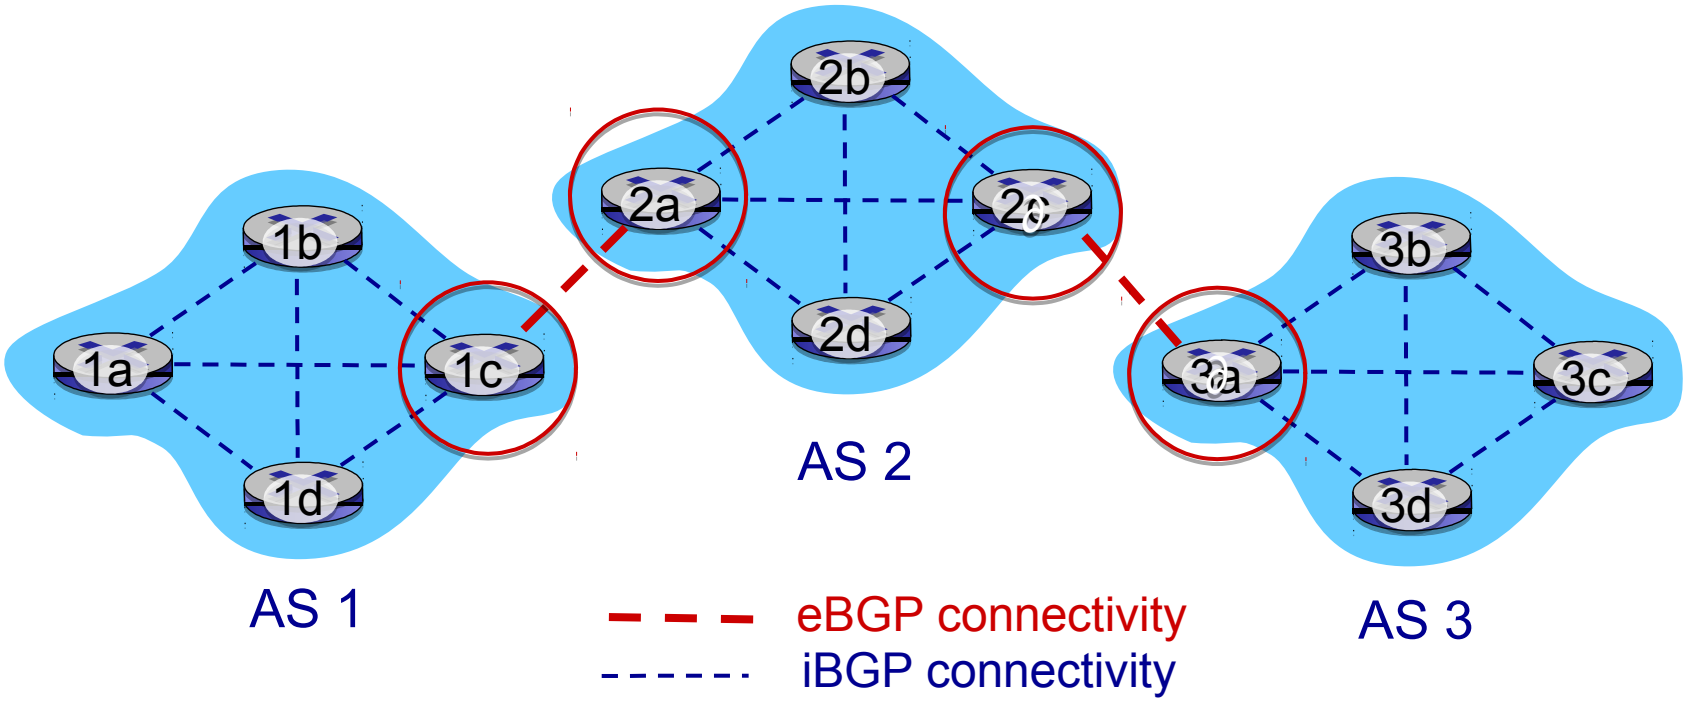

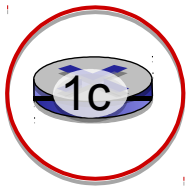

gateway routers run both eBGP and iBGP protools

BGP basics

- BGP session: two BGP routers ("peers") exchange BGP messages over semi-permanent TCP connection:
  - advertising *paths* to different destination network prefixes (BGP is a "path vector" protocol)
- when AS3 gateway router 3a advertises path AS3,X to AS2 gateway router 2c:
  - AS3 *promises* to AS2 it will forward datagrams towards X

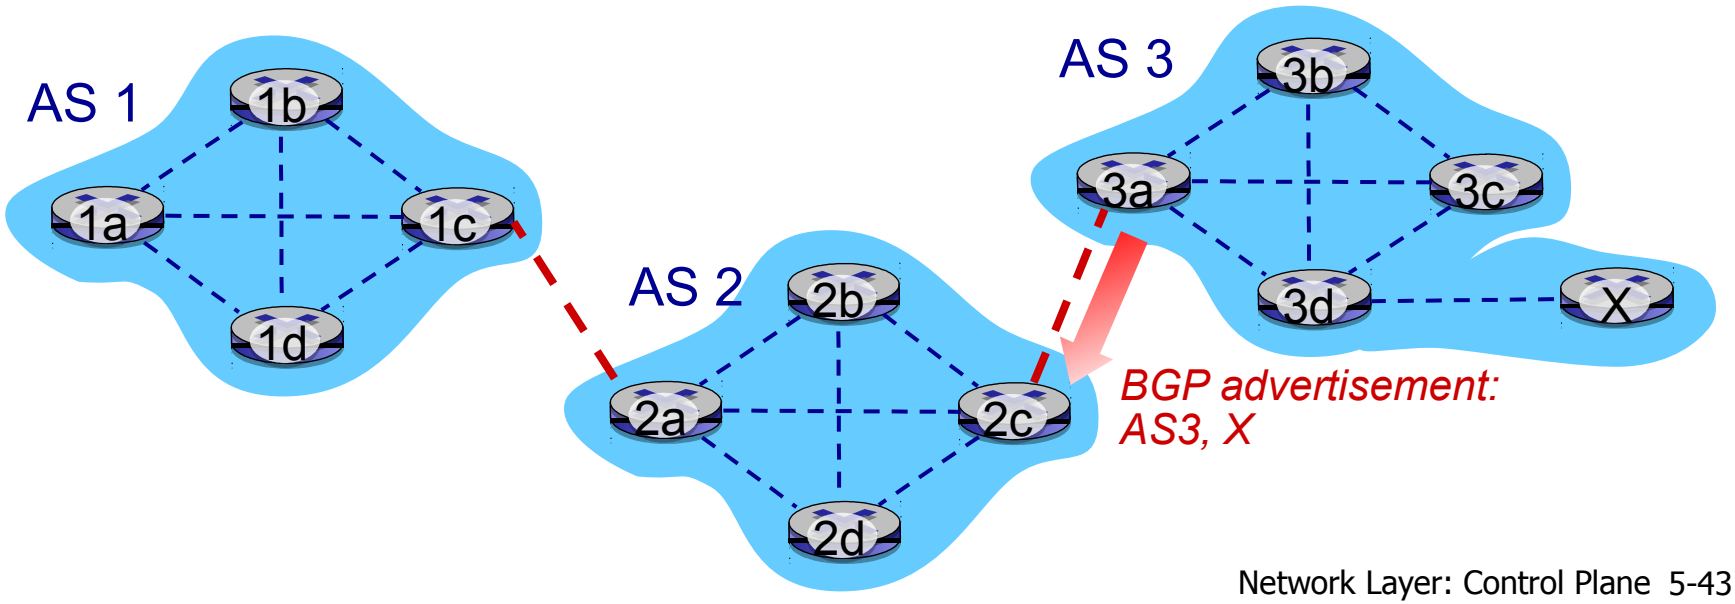

Path attributes and BGP routes

- advertised prefix includes BGP attributes
  - prefix + attributes = "route"
- two important attributes:
  - AS-PATH: list of ASes through which prefix advertisement has passed
  - NEXT-HOP: indicates specific internal-AS router to next-hop AS
- *Policy-based routing:*
  - gateway receiving route advertisement uses *import policy* to accept/decline path (e.g., never route through AS Y).
  - AS policy also determines whether to *advertise* path to other neighboring ASes

BGP path advertisement

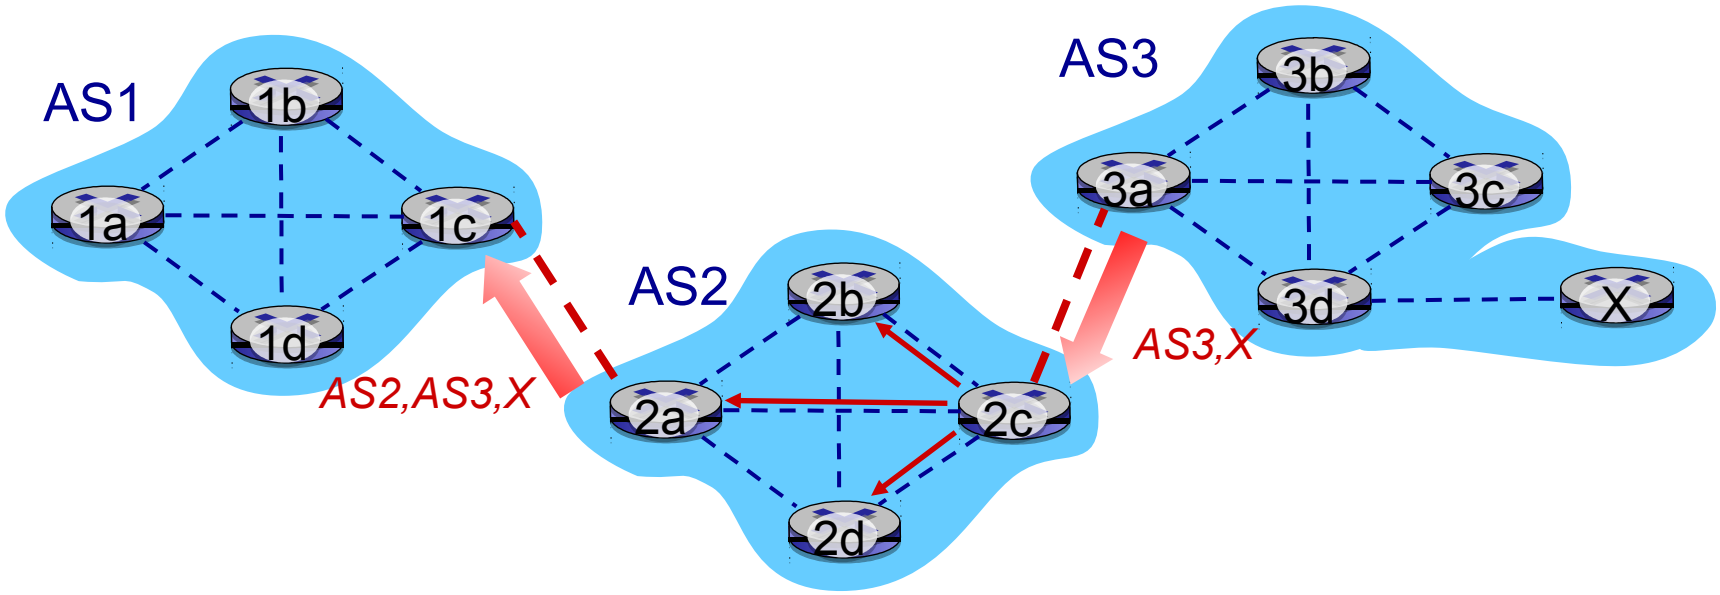

- AS2 router 2c receives path advertisement AS3,X (via eBGP) from AS3 router 3a
- Based on AS2 policy, AS2 router 2c accepts path AS3,X, propagates (via iBGP) to all AS2 routers
- Based on AS2 policy, AS2 router 2a advertises (via eBGP) path AS2, AS3, X to AS1 router 1c

BGP path advertisement

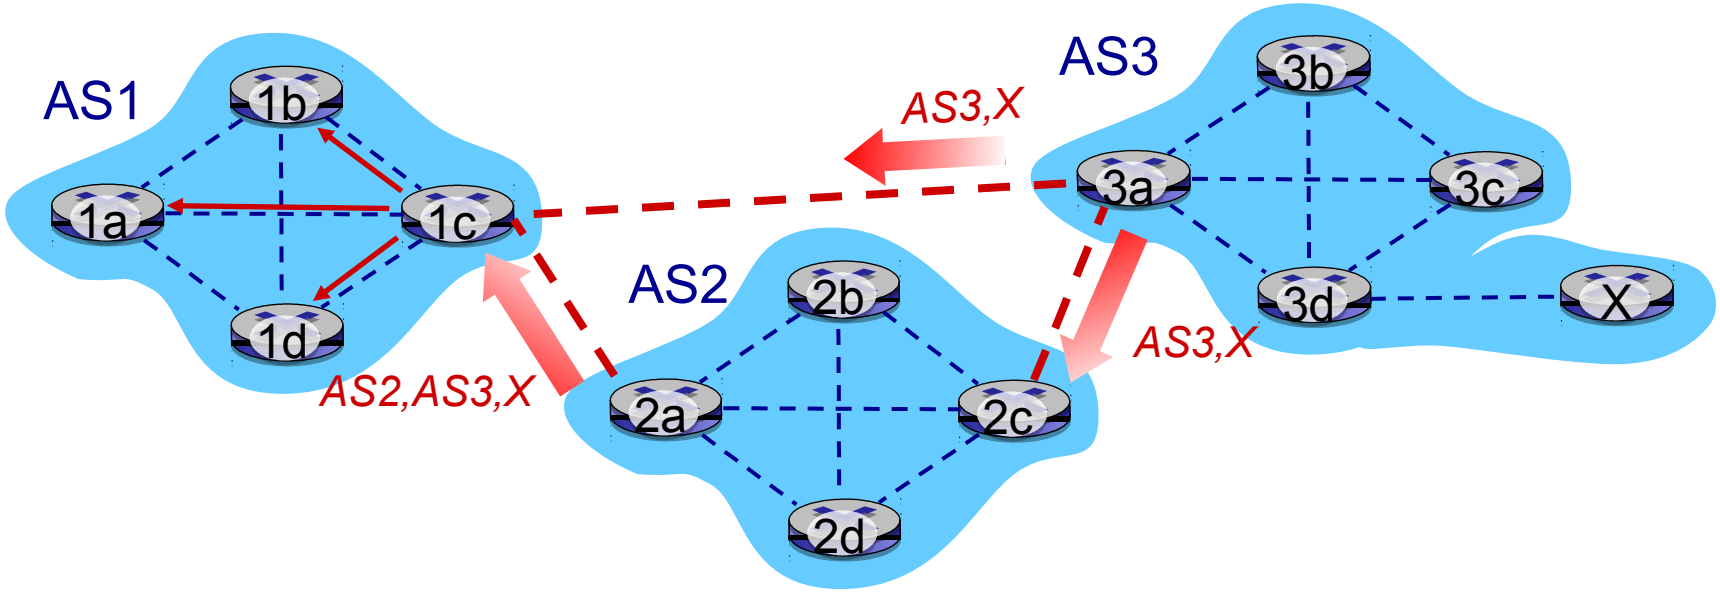

gateway router may learn about multiple paths to destination:

- AS1 gateway router 1c learns path *AS2,AS3,X* from 2a
- AS1 gateway router 1c learns path *AS3,X* from 3a
- Based on policy, AS1 gateway router 1c chooses path *AS3,X, and advertises path within AS1 via iBGP*

BGP messages

- BGP messages exchanged between peers over TCP connection
- BGP messages:
  - OPEN: opens TCP connection to remote BGP peer and authenticates sending BGP peer
  - UPDATE: advertises new path (or withdraws old)
  - KEEPALIVE: keeps connection alive in absence of UPDATES; also ACKs OPEN request
  - NOTIFICATION: reports errors in previous msg; also used to close connection

BGP: achieving policy via advertisements

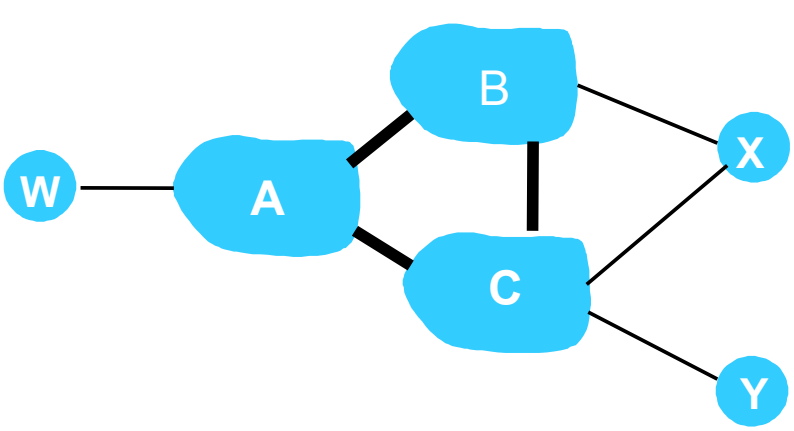

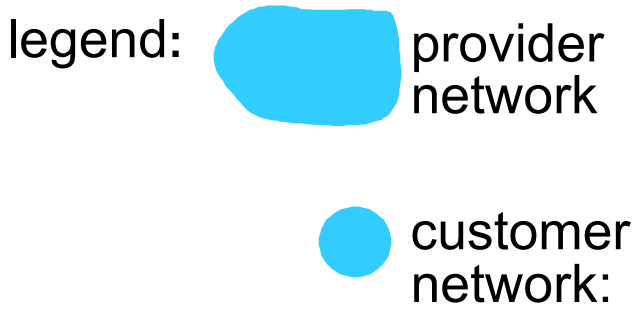

Suppose an ISP only wants to route traffic to/from its customer networks (does not want to carry transit traffic between other ISPs)

- A advertises path Aw to B and to C
- B *chooses not to advertise* BAw to C:
  - B gets no "revenue" for routing CBAw, since none of C, A, w are B's customers
  - C does not learn about CBAw path
- C will route CAw (not using B) to get to w

BGP: achieving policy via advertisements

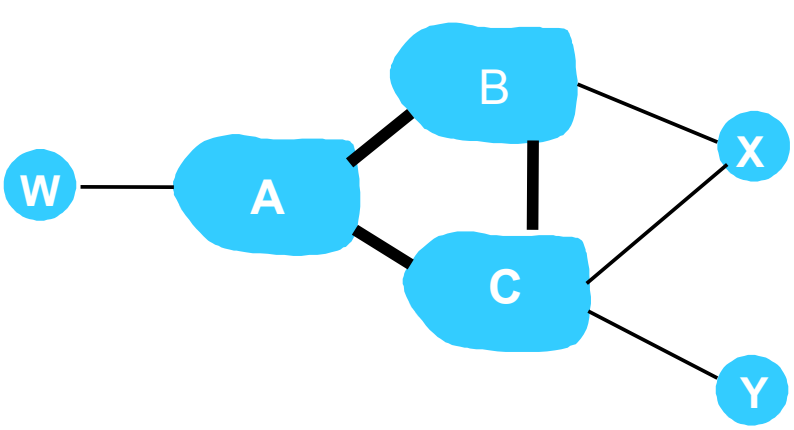

Suppose an ISP only wants to route traffic to/from its customer networks (does not want to carry transit traffic between other ISPs)

- A,B,C are *provider networks*
- X,W,Y are customer (of provider networks)
- X is *dual-homed:* attached to two networks
- *policy to enforce:* X does not want to route from B to C via X
  - .. so X will not advertise to B a route to C

Processamento de logs

- **Sistemas distribuídos nem sempre terão uma tela mostrando o que está acontecendo**
  - Mesmo que tenha → muita informação
  - Importante armazenar comportamento (anomalias, análise forense)
- **No passado: cada sistema (kernel e daemons principalmente) tinham seu próprio sistema de logs**
  - Reinventar a roda
  - Administrador tinha que aprender especificidades

Processamento de logs

- **Padronização de sistemas de logs**
  - syslog: protocolo que define como escrever logs (*facility*, *severity* e mensagem)
  - Pode ser implementado em um servidor fisicamente separado dos servidores onde os daemons a serem monitorados estão rodando
  - Aceita conexões remotas (UDP porta 514, TCP porta 6514)
- **Implementações**
  - rsyslog
  - syslog-ng

Processamento de logs

- **O que registrar?**
  - Instante de tempo
  - Importância do evento
  - Descrição do evento
- **Cuidados**
  - Relógios sincronizados (protocolo NTP é essencial)
  - Cuidado com logs “verbose”
    - /var/log geralmente em discos com replicação e backup constante

Processamento de logs

- **Como processar os logs depois?**
  - Considerar como processamento de fluxo
  - Importante normalizar
    - logstash é uma ferramenta muito usada
    - PLN em alguns casos (tudo depende de como o log foi gerado)
      - » Importante registrar as mensagens lembrando que provavelmente uma máquina vai processar, não um ser humano

Referências

- J. F. Kurose, K. W. Ross. Computer Networking, A Top-Down Approach (7th edition), Pearson, 2016;
- Wikipedia. “syslog”. https://en.wikipedia.org/wiki/Syslog
- Logstash: Colete, analise e transforme logs | Elastic. https://www.elastic.co/pt/logstash