CSCD 433/533
Network Programming
Winter 2017

Lecture 12
Global Address Space
Autonomous Systems, BGP Protocol
Routing

Chapter 5

Topics

Interdomain Routing

- BGP – Interdomain Routing
- Benefits vs. Link State Routing
- BGP Operation
- Interior vs Exterior Routing
- Limitations of BGP

Connected to Internet

- What does it mean to be connected to the Internet?
  - Packets sent to host arrive at host
  - Packets sent back arrive at destination

- Means must have a path to you
- Your ISP must have a path to you
- My IP must lie within an address space that gets advertised as a route by others

- Else no-one can find me

BGP's Role in Connectivity

Promises made. Promises kept.

- Advertises routes
  - Think of these as promises to carry data to IP space represented in route being advertised
- You promise that If someone sends data to route that is advertised
  - Know how to carry data to its ultimate destination
  - Don't want to advertise routes to places that we don't know how to reach !!!

Routing Inside and Outside AS's

- Internet organized into AS's
  - Routing is divided into between AS's and routing within AS's
  - Creates another level of hierarchy

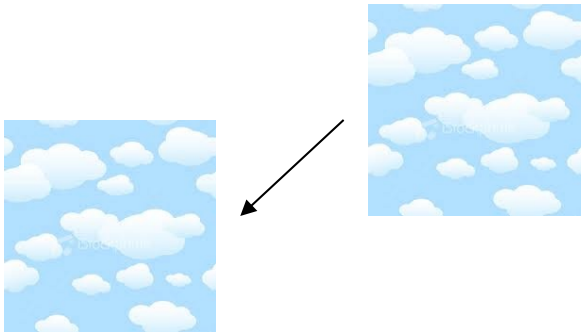

  - Today's Internet
    - Number of backbone networks operated by private companies
    - Smaller ISP's have service agreements with larger ISP's
    - Some only provide service to end users

Routing Between and Within AS's

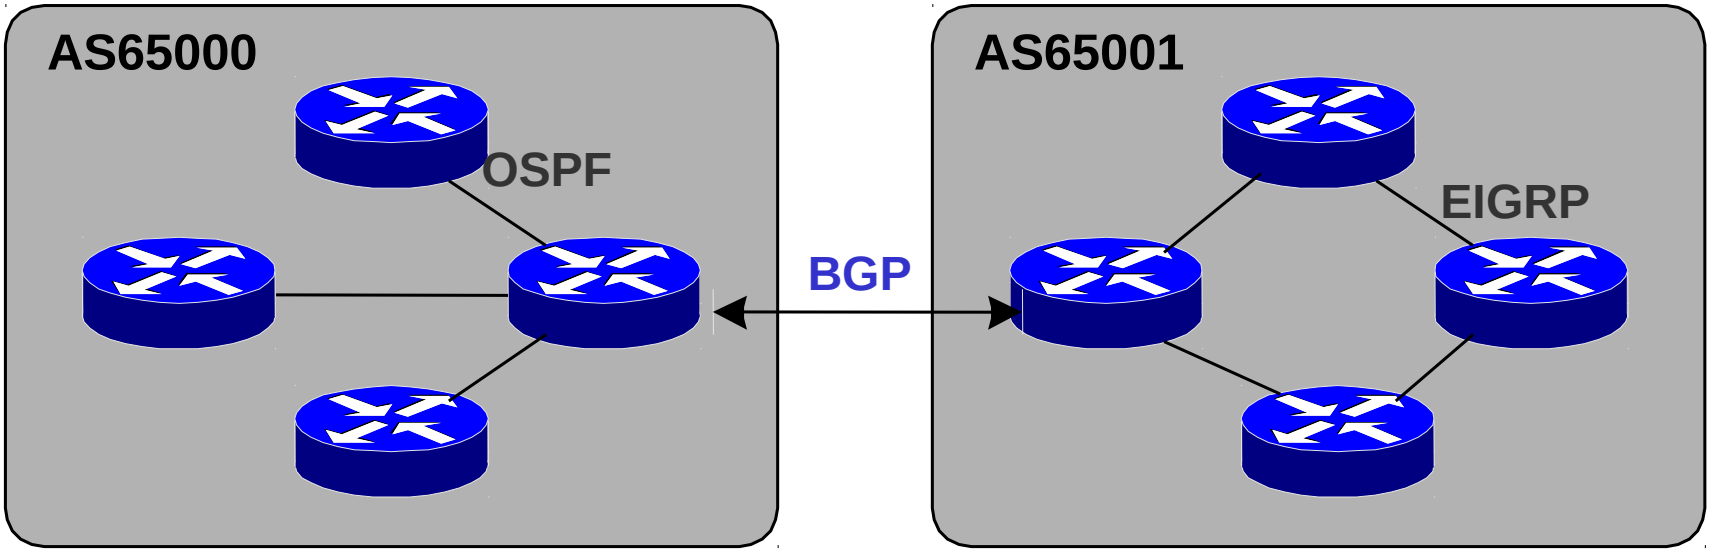

Interior routing protocol (IGP) runs inside an autonomous system resulting in optimum intra-AS routing

- Exterior Gateway Protocol (EGP) run between autonomous systems to enable routing policies

Interior Gateway Protocol IGP

- **Within** a network/autonomous system
- Carries information about internal infrastructure prefixes

- Examples – OSPF, ISIS, EIGRP, RIP and IGP

Exterior Gateway Protocol EGP

- Used to convey routing information between networks/ASes
- De-coupled from the IGP
- Current EGP is BGP4

EGP is Currently BGP

- Why do we need a separate routing protocol for interAS routing?
  - Many reasons ... but essentially

Purpose of Exterior routing is different!!!

- AS's route based on business relationships
- Not really about optimizing routes
- Need a protocol to recognize relationships
- Change routing tables and not disrupt internal routing

Shortest-Path Routing is Restrictive

- All traffic must travel on shortest paths
- All nodes need common notion of link costs
- Incompatible with commercial relationships

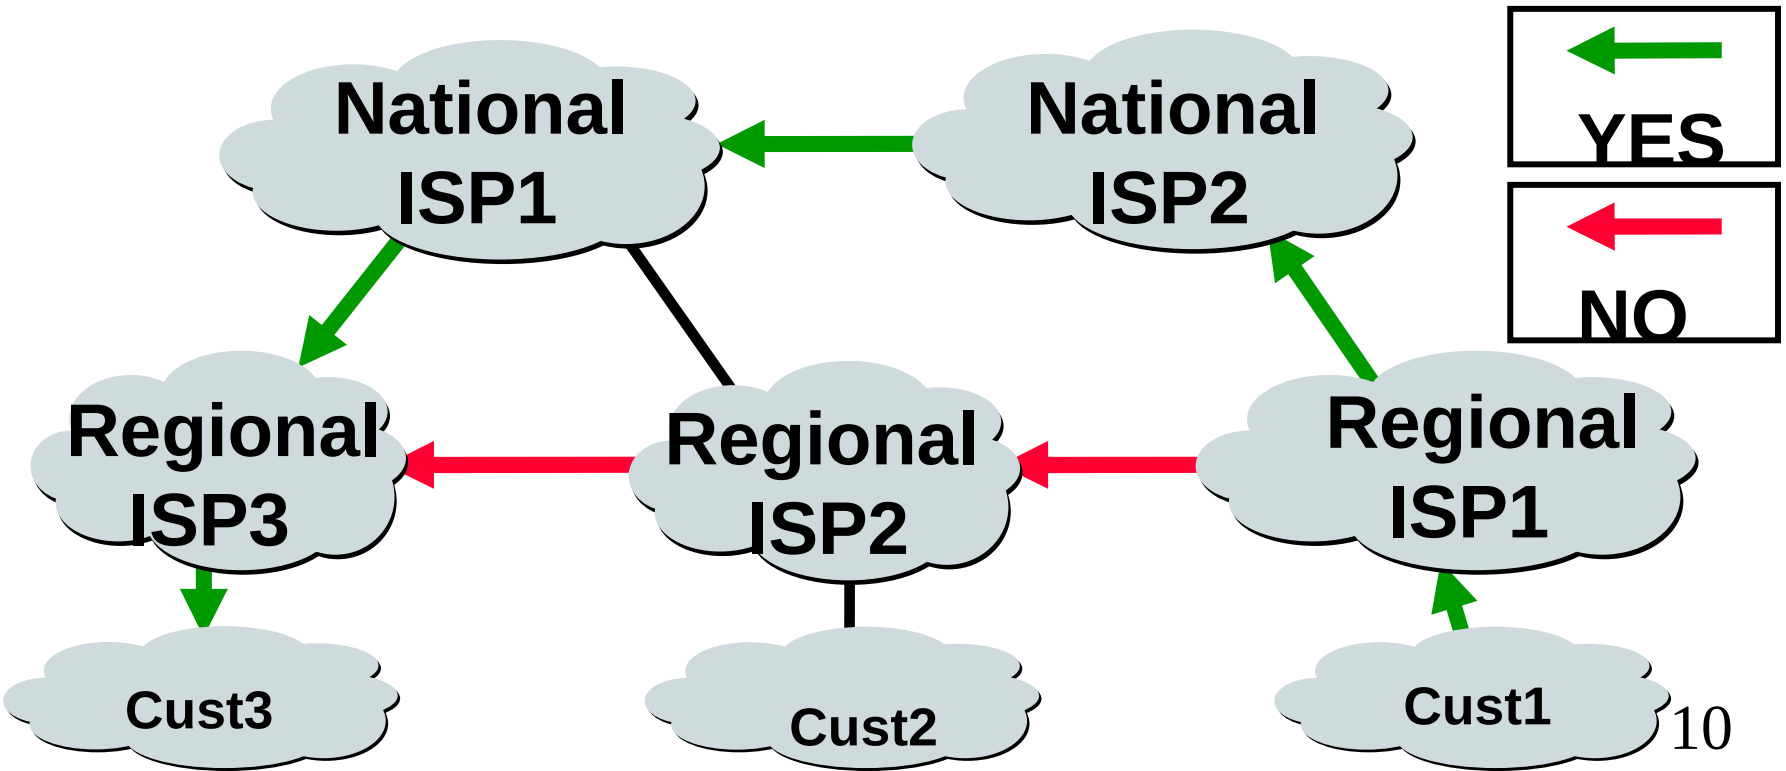

More Problems with Link-State Shortest Path Routing

- Topology information is flooded
  - High bandwidth and storage overhead
  - Forces nodes to divulge sensitive information
- Entire path computed locally per node
  - High overhead for large network
- Minimizes some notion of total distance
  - Works only if policy is shared and uniform
  - Not true for Internet!!!

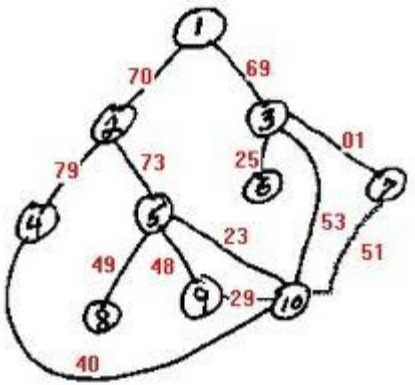

- Path Vector to the rescue ....

BGP is a Path Vector Protocol!!!

Path-Vector Routing

- Extension of distance-vector routing (RIP)
  - Supports flexible routing policies
  - Avoids count-to-infinity problem
- Key idea: Advertise Entire Path
  - Distance Vector: Send distance metric per dest d
  - Path Vector: Send entire path for each dest d

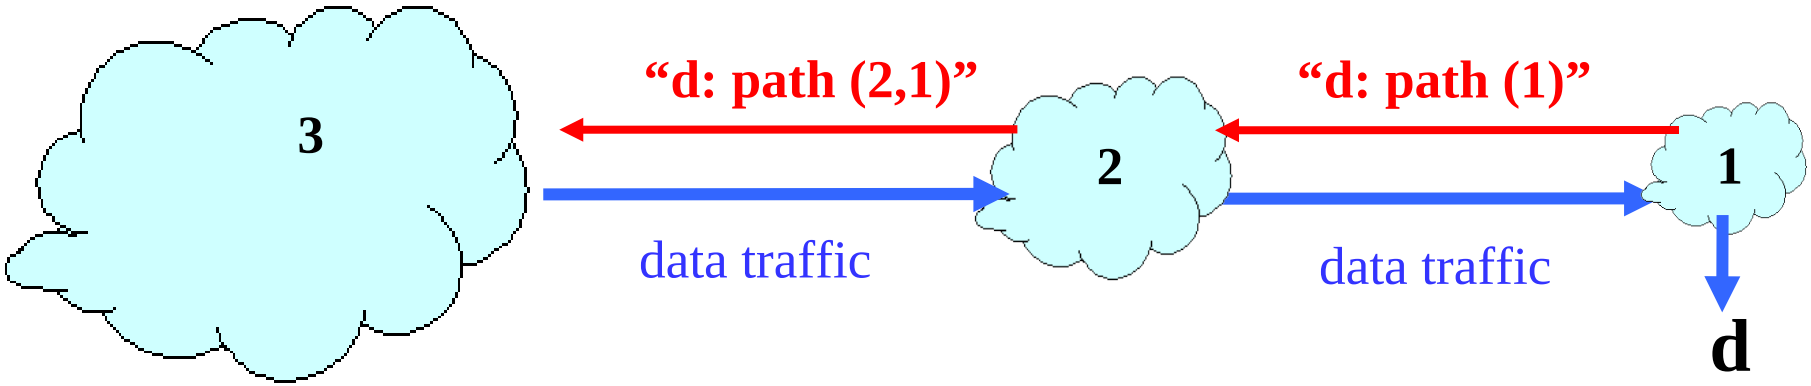

Faster Loop Detection

- Node can easily detect a loop
  - Look for its own node identifier in path
  - E.g., node 1 sees itself in the path “3, 2, 1”
- Node can simply discard paths with loops
  - E.g., node 1 simply discards the advertisement

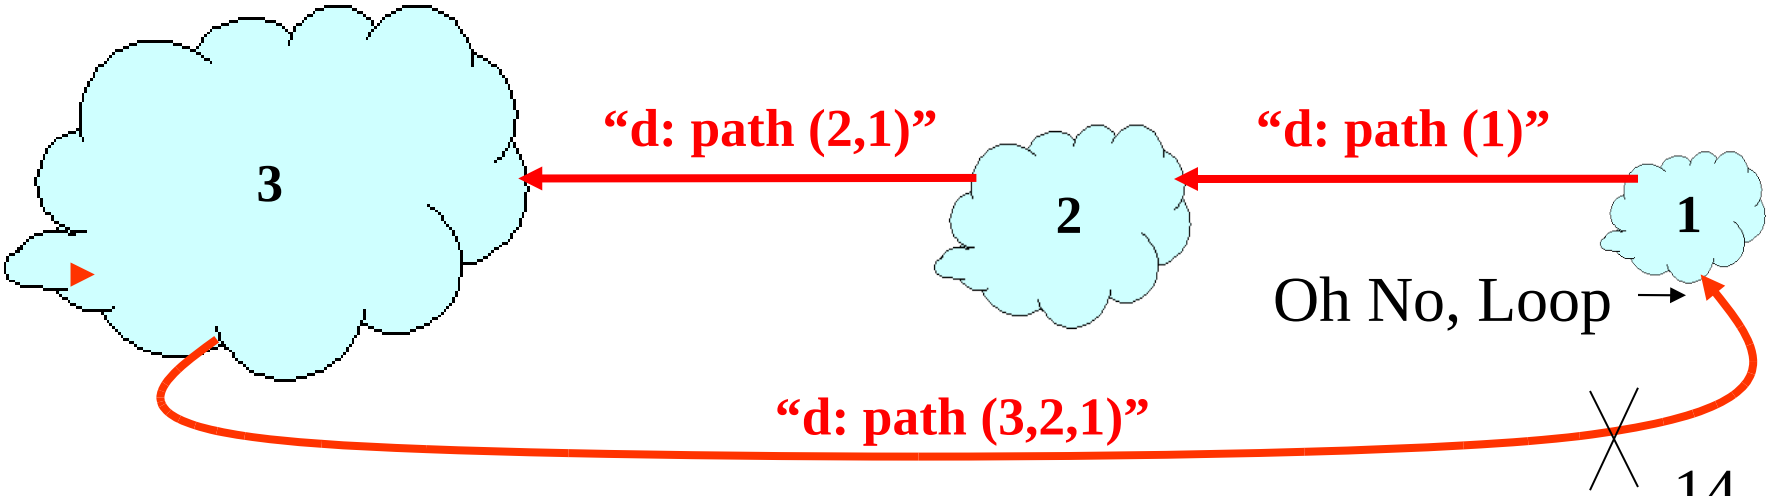

Flexible Policies

- Each node can apply local policies
  - Path Selection: Which path to use?
  - Path Export: Which paths to advertise?
- Examples
  - Node 2 may prefer the path “2, 3, 1” over “2, 1”
  - Node 1 may not let node 3 hear the path “1, 2”

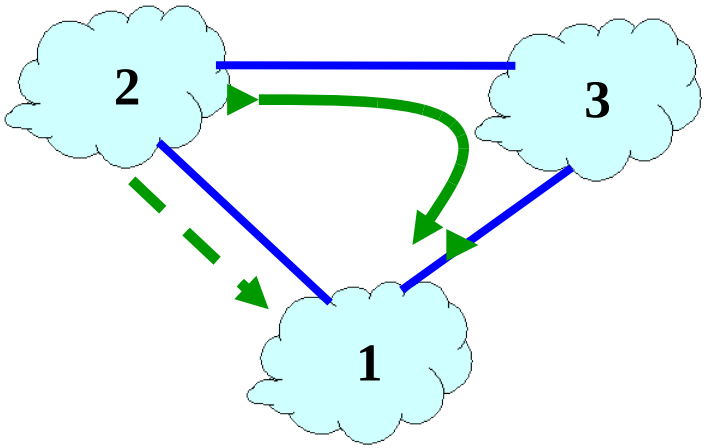

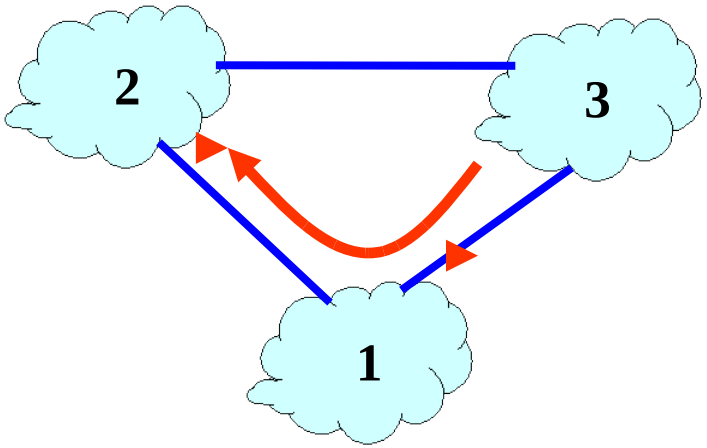

Internet Backbone and Relationships

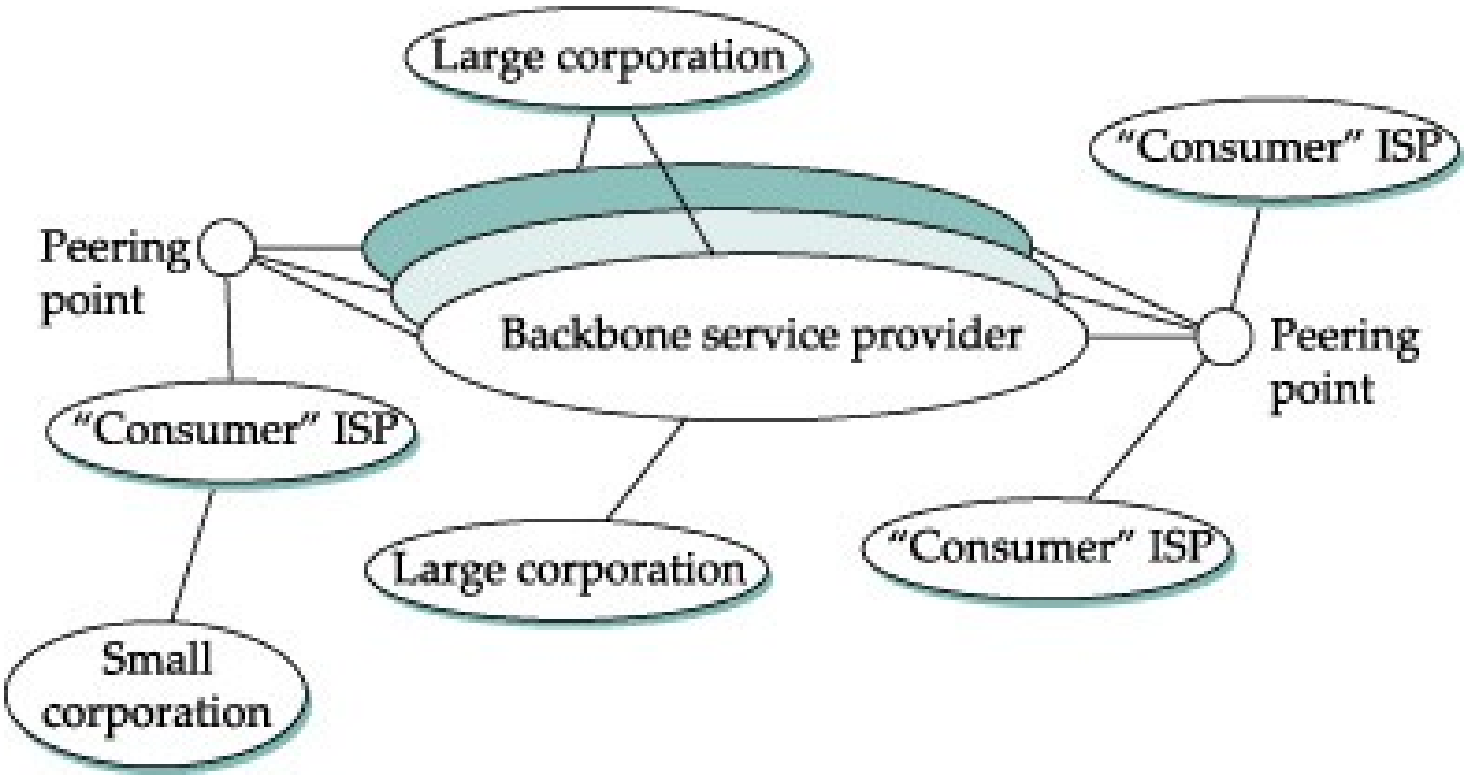

AS's Have Business Relationships

- Neighboring AS's have business contracts with each other
  - How much traffic to carry?
  - Which destinations to reach
  - How much money to pay for transport of traffic
- Common business relationships
  - Customer-transit
    - AT&T allows Sprint to to use their routes
  - Customer-provider
    - E.g., EWU is a customer of AT&T, the provider
  - Peer-peer (think of them as equal)
    - E.g., AT&T is a peer of Sprint

AS's Have Business Relationships

- Peer-Peer (equal)
- Voluntary connection between administratively separate networks
- Exchange traffic between customers of each network
- Requires physical interconnection of networks
- Have peering agreements from "handshake" to thick contracts

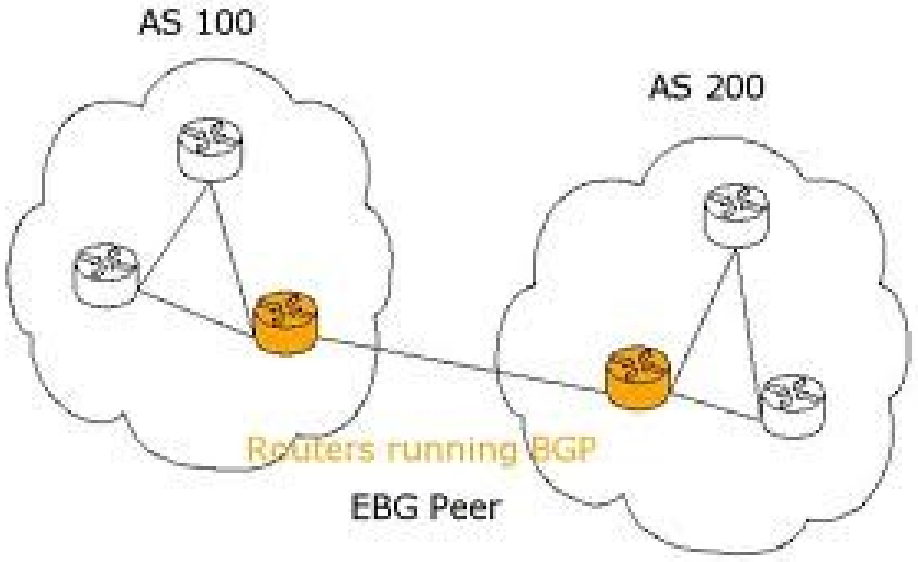

AS's Have Business Relationships

Transit

- An ISP pays money to another for Internet access via them

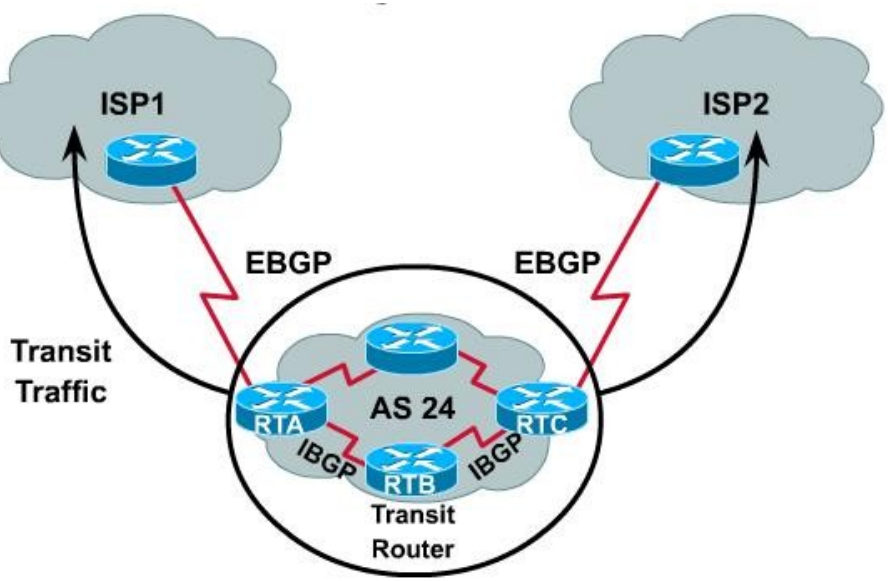

Customer

- Customer of ISP pays them to advertise their route

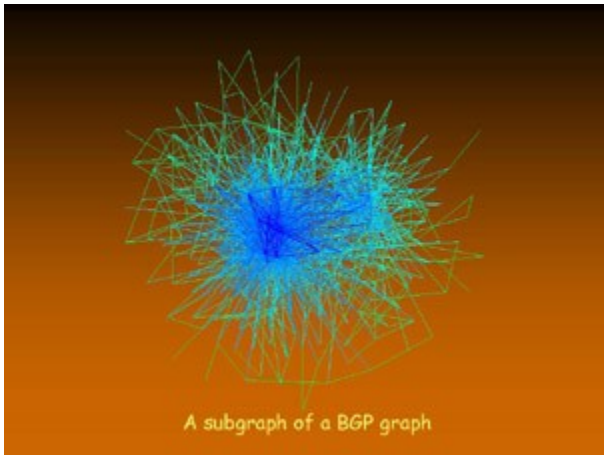

BGP Border Gateway Protocol

History of BGP

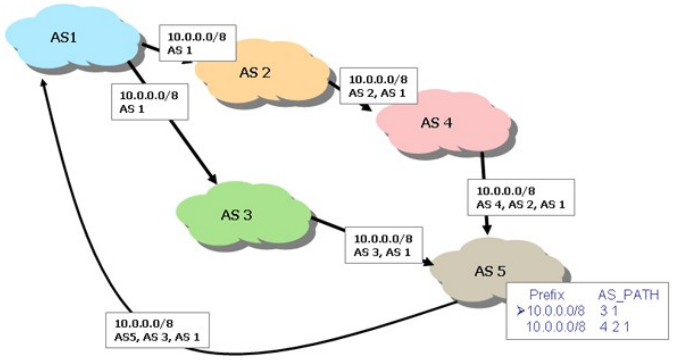

- BGP became Internet standard **1989**
  Originally defined in RFC 1105
- Current version, BGP-4, was adopted in **1995**
  and is defined in RFC 1771 about 17 years old
- BGP-4 supports Classless Inter Domain Routing
  (CIDR)
- Recently updated, 2006 – RFC 4271

IS the only routing protocol used today to route between autonomous systems

Who Can Run BGP?

- If you want to run BGP,
  - Ask ISP to see if they will agree to communicate with you via BGP
  - You will have to show your need to run BGP
  - Routes that run BGP are T1 speeds and above
  - Only providers that allow you to exchange BGP routes with them are Internet Service Providers

Border Gateway Protocol (BGP)

- Maintains table of IP networks or 'prefixes' which designate network reachability among Autonomous Systems (AS)
  - Described as a path vector protocol
    - Instead of hop count, uses AS's as hops
  - BGP makes routing decisions based
    - On path,
    - Network policies and/or
    - Rulesets

Border Gateway Protocol (BGP)

- No metrics
- Not about optimizing anything
- All about policy (business and politics) …

Politics and Business

BGP Has Two Versions

- **Two Versions of BGP**
  - Routers between AS's use EBGP
  - Routers within AS can use iBGP to synchronize tables, BGP speaking routers use iBGP

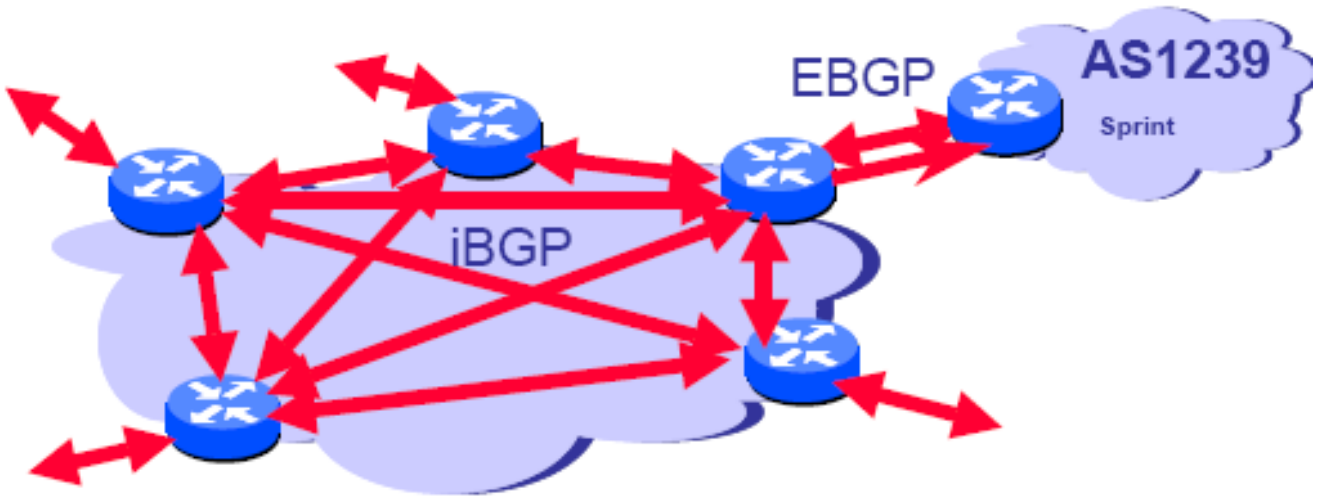

BGP Has Two Versions

- BGP uses the same types of message on IBGP and EBGP sessions
  - But rules for when to send which message and how to interpret each message differ slightly
    - For this reason
    - Some people refer to IBGP and EBGP as two separate protocols

Border Gateway Protocol (BGP)

- Border routers in each AS communicate with neighboring routers in other AS's
- BGP route announcements say
  - "I can reach this network, and this is the path of AS Numbers I heard this from
  - Plus some attributes I choose to tell you
  - Can't accept route if your AS Number is in it"

Border Gateway Protocol

- BGP works by advertising a complete path of AS's to reach a particular network
- Example network … Details follow
  - How to get to 128.96, 192.4.153 etc.?

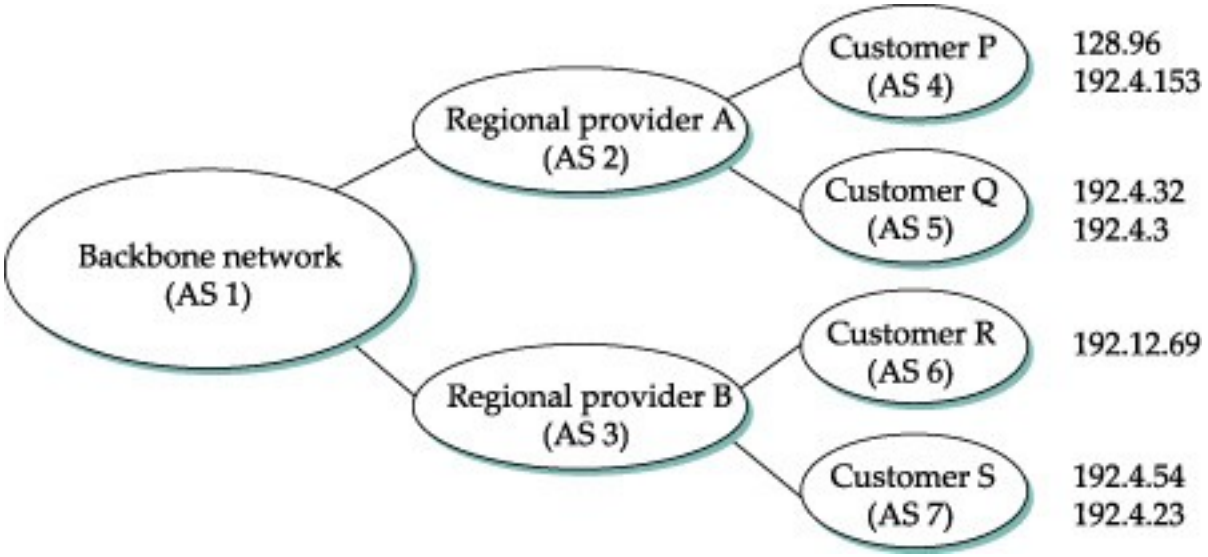

BGP Protocol

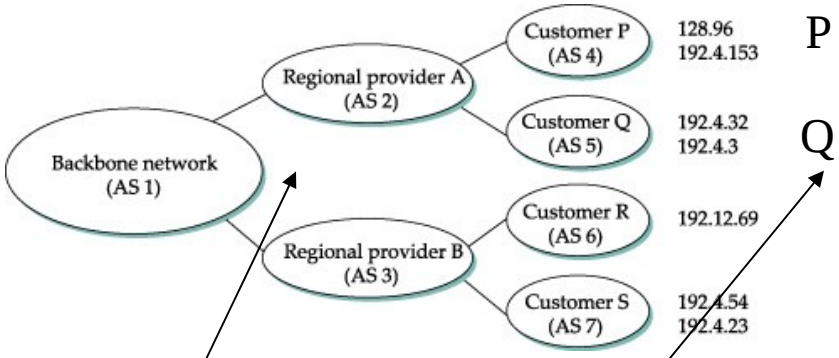

- Example: Details
- BGP works by electing at least one Speaker Router for the entire AS of Provider A (AS2)
  - Able to advertise reachability info for each network assigned to customers P and Q
  - **Thus AS2 would say**
    - Network 128.96, 192.4.153, 192.4.32 and 192.4.3 can be reached directly from AS2
    - Backbone network, AS1 advertises 128.96, 192.4.153, 192.4.32, 192.4.3 can be reached along path (AS1, AS2)

More BGP Details

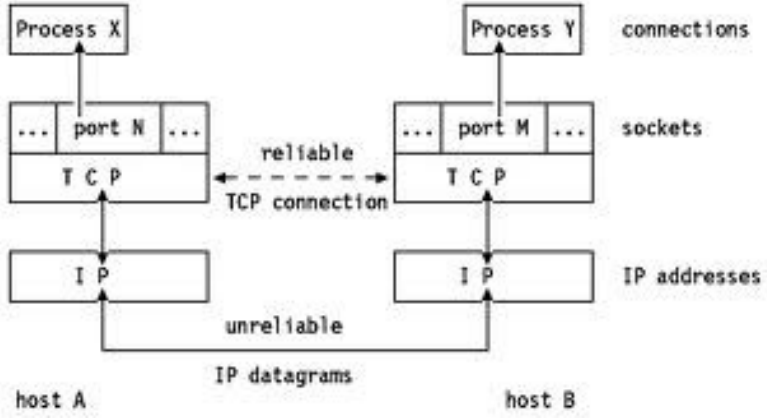

- Uses TCP as its Transport Protocol
  - Uses port **179**
  - Use of **TCP** as BGP's transport protocol
    - Guarantees reliability
    - Eliminates additional complexity related designing reliability into protocol itself
  - **BGP data enclosed within TCP packets**
    - After setting up BGP session and exchanging initial routes,
    - BGP peers trade incremental routing and notification updates **What layer is BGP?**
    - Answer: 7

More BGP Details

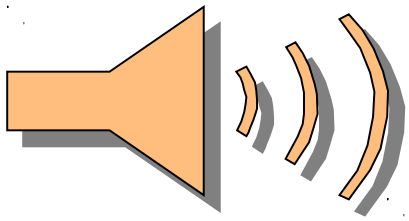

- Routers that run BGP routing process referred to as **BGP speakers**
  - Pair of BGP-speaking routers form TCP connection and exchange routing information
    - Become BGP neighbors
    - Also called “peers” have peering sessions
    - A single router can participate in many peering sessions at any given time
  - See next slide ...

BGP Operations (Simplified)

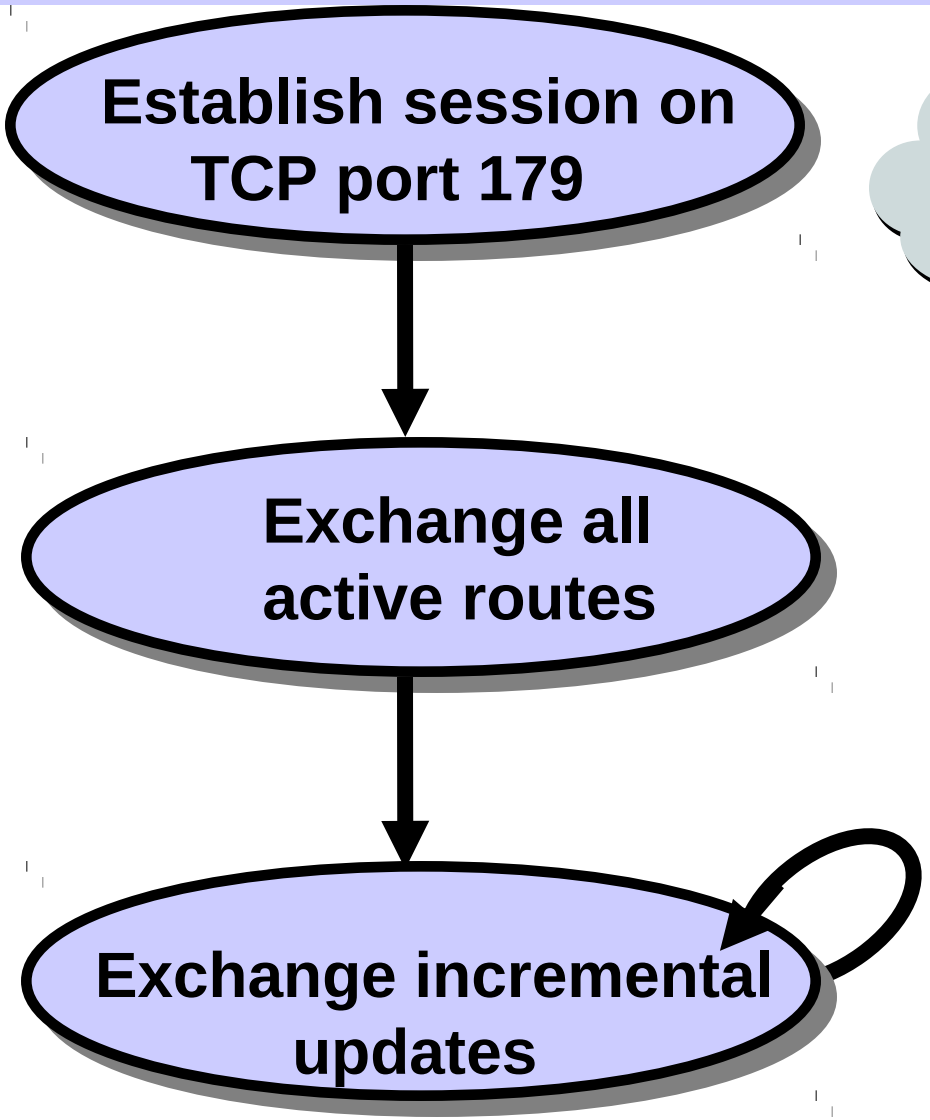

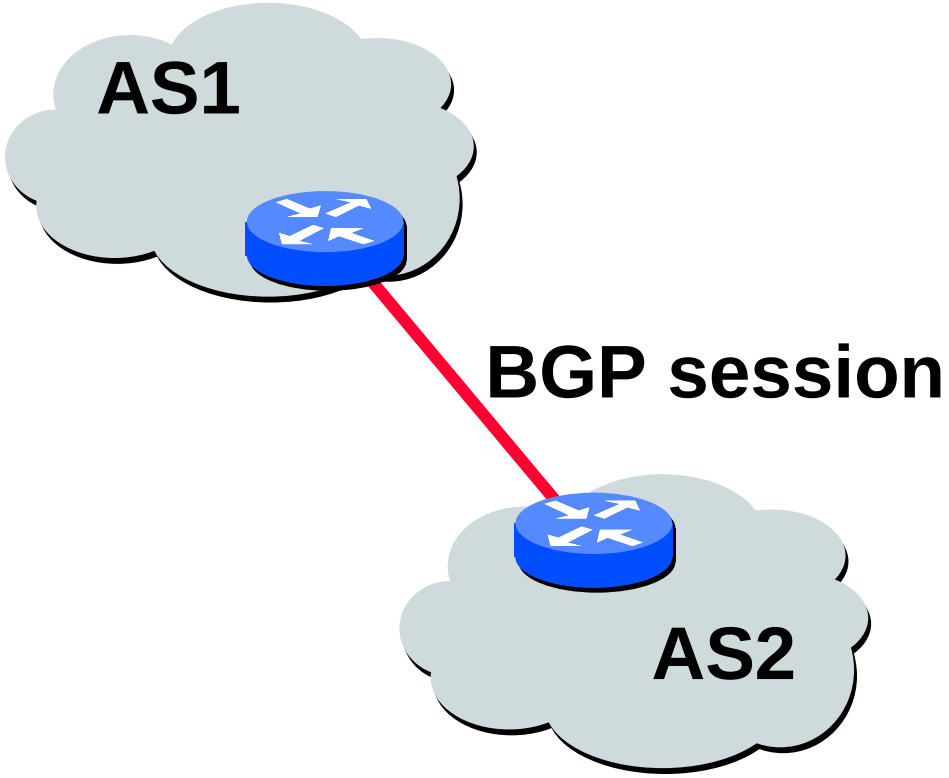

While connection is ALIVE exchange route UPDATE messages

BGP Sessions

- One router can participate in many BGP sessions
- Initially … node advertises ALL routes it wants neighbor to know about
- Ongoing … only inform neighbor of changes

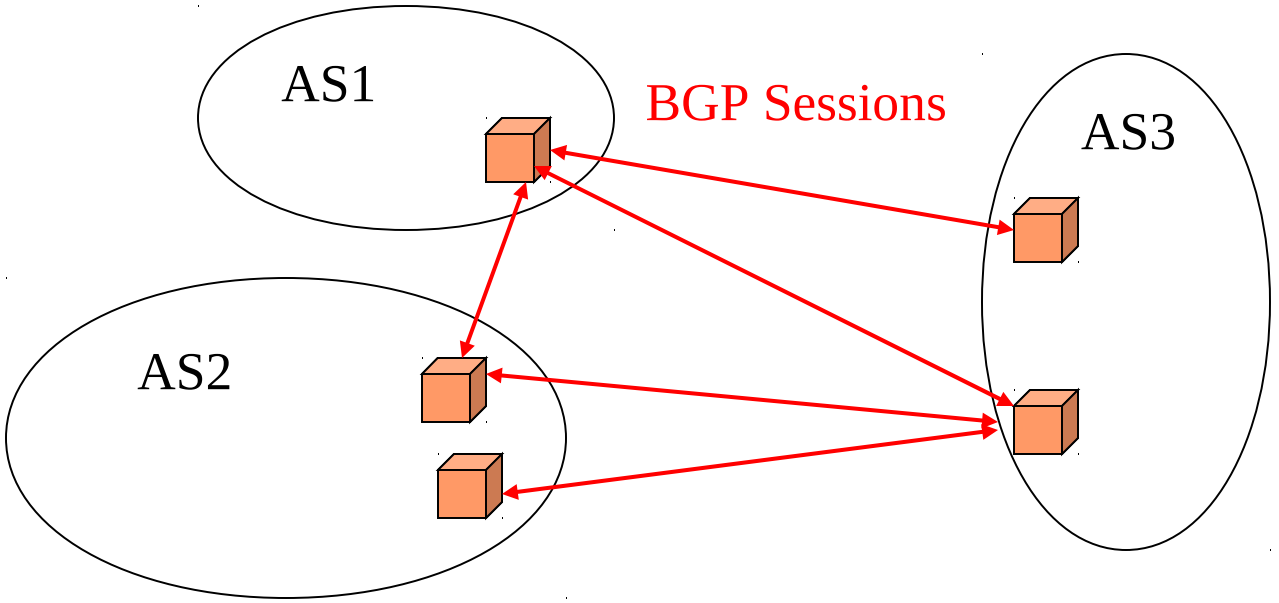

Four Basic Messages

- **Open**
  Establishes BGP session (TCP port #179)
- **Notification**
  Report unusual conditions
- **Update**
  Inform neighbor of new routes that become active
  Inform neighbor of old routes that become inactive
- **Keepalive**
  Inform neighbor that connection is still viable

OPEN Message

- During session establishment, two BGP speakers exchange their
  - AS numbers
  - BGP identifiers (usually one of the router's IP addresses)
- A BGP speaker has option to refuse a session
- Select value of a **hold timer**
  - Maximum time to wait to hear something from other end before assuming session is down
- Authentication information (optional)

Pictorial Diagram for BGP Finate State Machine

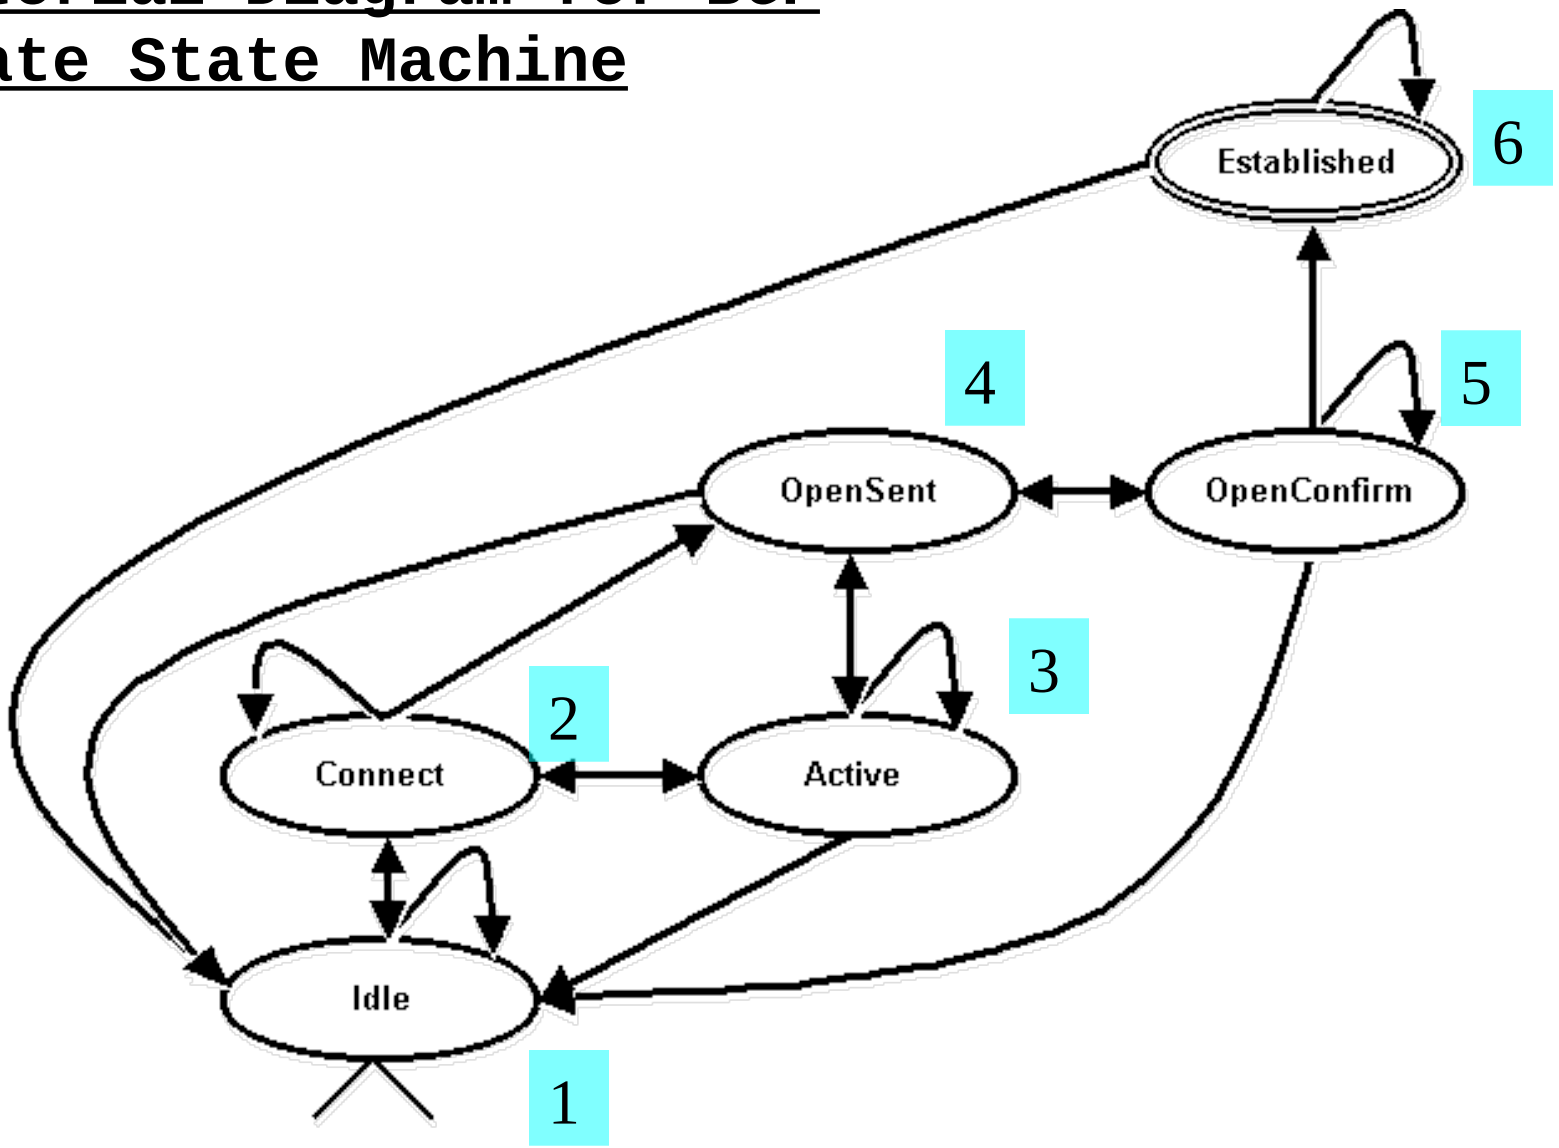

1. Idle

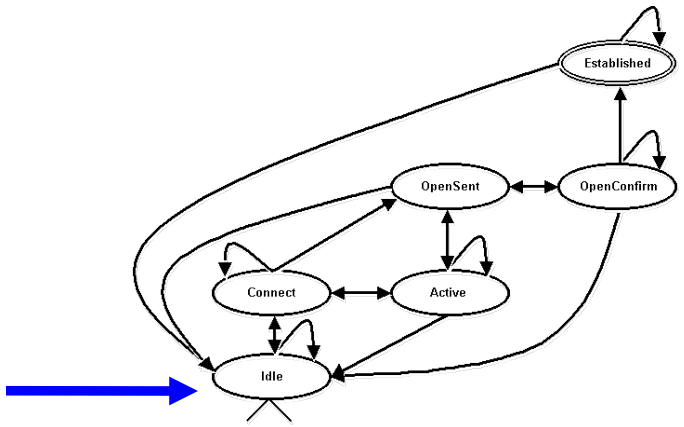

- BGP speaker router waits for session, in IDLE state
- Will not start session until **start event** occurs
- Initial configuration Open, or clearing of BGP session as start event and system transitions to connect state
- When BGP session shuts down because of error, it returns to Idle state
  - NOTIFICATION messages used to signal connection errors return router to this state

BGP FSM Explained

2. CONNECT

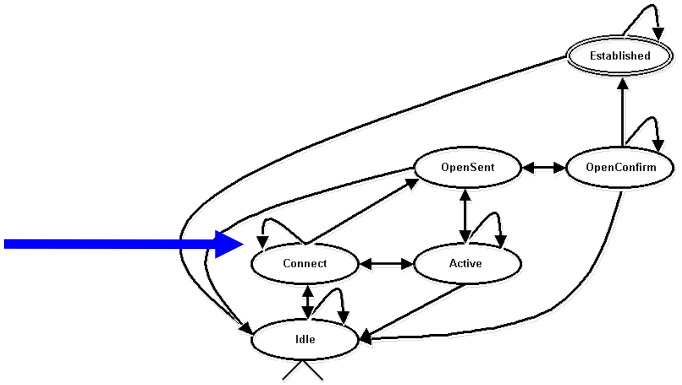

- Once BGP software and it's environment have been initialized,
- OPEN message is sent
- Router has attempted to open TCP connection between itself and another BGP speaking peer

BGP FSM Explained

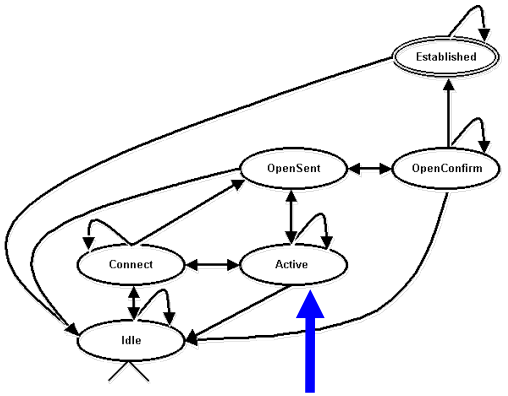

3. ACTIVE

- Router started first phase of initializing TCP three-way handshake to remote router (peer)
- If router fails to establish TCP connection, it drops back to **IDLE**.

BGP FSM Explained

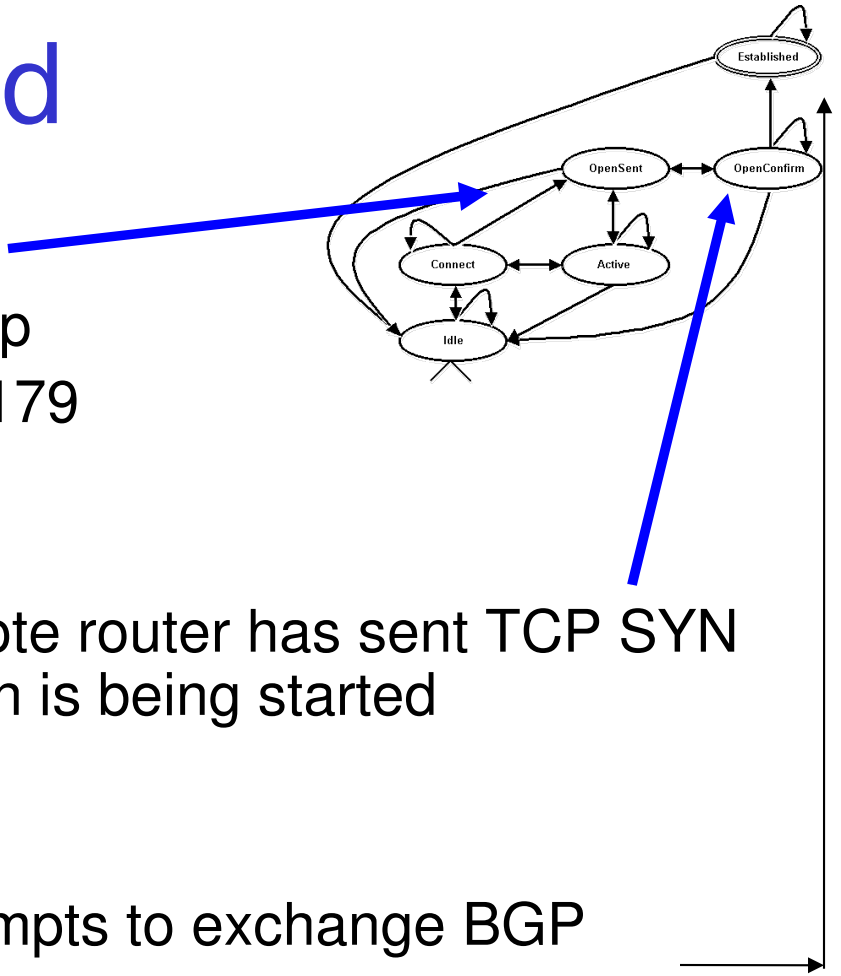

4. OPEN SENT

Once BGP has performed all setup steps, it sends TCP SYN on port 179

5. OPEN CONFIRM

Router enters this state after remote router has sent TCP SYN packet indicating that TCP session is being started

6. ESTABLISHED

After TCP handshake, router attempts to exchange BGP messages..

If router is in OPEN CONFIRM state and receives UPDATE or KEEPALIVE message, BGP session state changes to ESTABLISHED

BGP FSM Explained

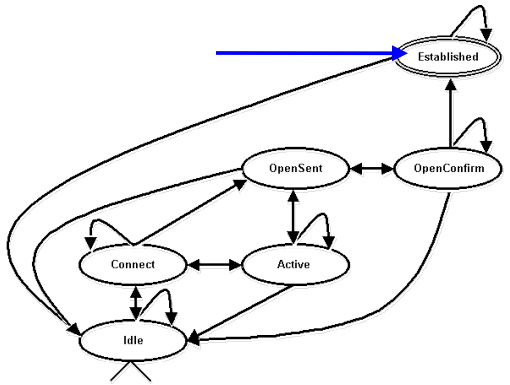

ESTABLISHED STATE

- This is the ONLY state in which BGP will actually exchange routes!!
- Established is ONLY state that counts
- Any other state, you have a non-functional BGP session (and possibly a broken physical link if it refuses to establish the connection)
- Stays in this state
  - Sends either Update or Keepalive messages

Review of Supernetting from CSCD330

http://www.2000trainers.com/cisco-ccna-05/ccna-classless-cidr-supernetting/

- Want to aggregate 8 network addresses between 131.0.0.0/16 and 131.7.0.0 /16
- So, range can now be designated as 131.0.0.0/13 This value aggregates all addresses between 131.0.0.1 and 131.7.255.254

/13 Subnet Mask	11111111	11111	000	00000000	00000000	
	Network					
131.0.0.0	10000011	00000	000	00000000	00000000	
131.1.0.0	10000011	00000	001	00000000	00000000	
131.2.0.0	10000011	00000	010	00000000	00000000	
131.3.0.0	10000011	00000	011	00000000	00000000	
131.4.0.0	10000011	00000	100	00000000	00000000	
131.5.0.0	10000011	00000	101	00000000	00000000	
131.6.0.0	10000011	00000	110	00000000	00000000	
131.7.0.0	10000011	00000	111	00000000	00000000	
Network ID	10000011	00000	000	00000000	00000000	= 131.0.0.0/13

BGP Routing Details

- BGP rules state that longer routes are more specific and preferred, more bits for network portion
- So, for example, YouTube, owns IP space
  - 208.65.153.0/24,
  - 208.65.152.0/24 and
  - 208.65.154.0/23,
- YouTube announces single aggregated BGP route for /24 prefixes, announced as 208.65.152.0/22

208.65.152.0/22 via AS 36561 (YouTube)

BGP Enables Policy Based Routing

- BGP provides mechanisms for policy-based routing
  - BGP routers can
    - **Rank** routes and **Control** information redistribution
  - BGP carries out policy routing by
    - Filtering routes, based on attributes
- Policies are not part of the protocol!!!
  - Decisions made by AS administrator, specified to BGP in configuration files
  - Routing policies often based on
    - Security,
    - Economic, or
    - Political considerations

Attributes Used to Select Best Routes

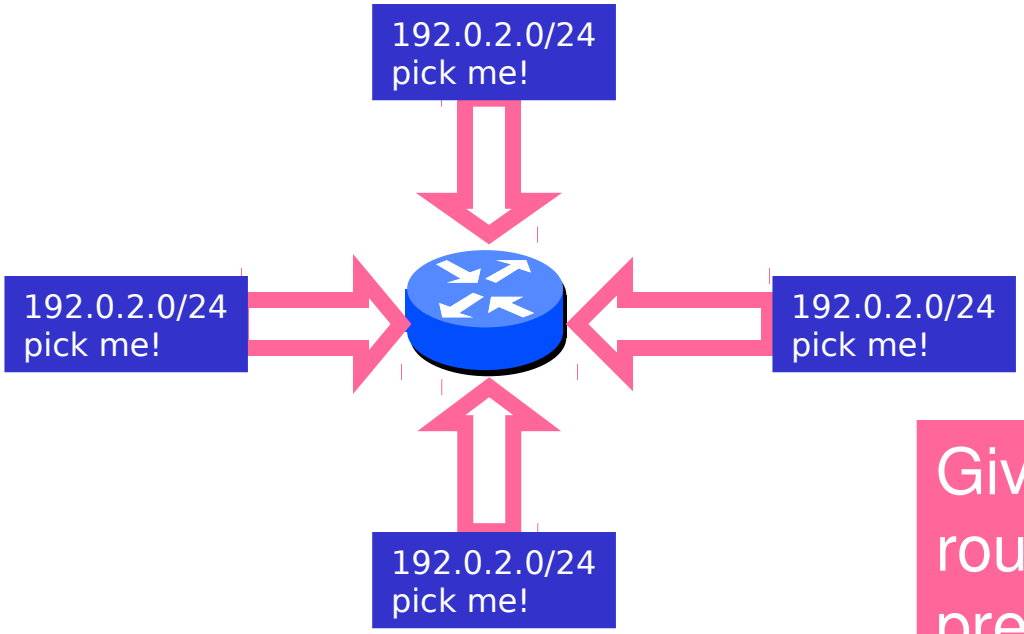

Look at these ...

Given multiple routes to the same prefix, a BGP speaker must pick at most one best route

(Note: it could reject them all!)

Policy Attributes - Mandatory

- **ORIGIN**
  ORIGIN is mandatory attribute that indicates origin of prefix, or rather, way in which prefix was injected into BGP
  There are three origin codes, listed in order or preference:
  - **IGP**, says prefix originated from information learned from an interior gateway protocol
  - **EGP**, says prefix originated from EGP protocol, which BGP replaced
  - **INCOMPLETE**, says prefix originated from some unknown source

Policy Attributes - Mandatory

- AS-PATH
  - List of AS's through which announcement for prefix has passed
  - Each AS prepends its AS # to AS-PATH attribute when forwarding an announcement
  - Useful to detect and prevent loops
  - Shorter AS-Path is preferred

Prefix	***Next hop***	***AS Path***
128.73.4.21/21	232.14.63.4	1239 701 3985 631

ASPATH Attribute

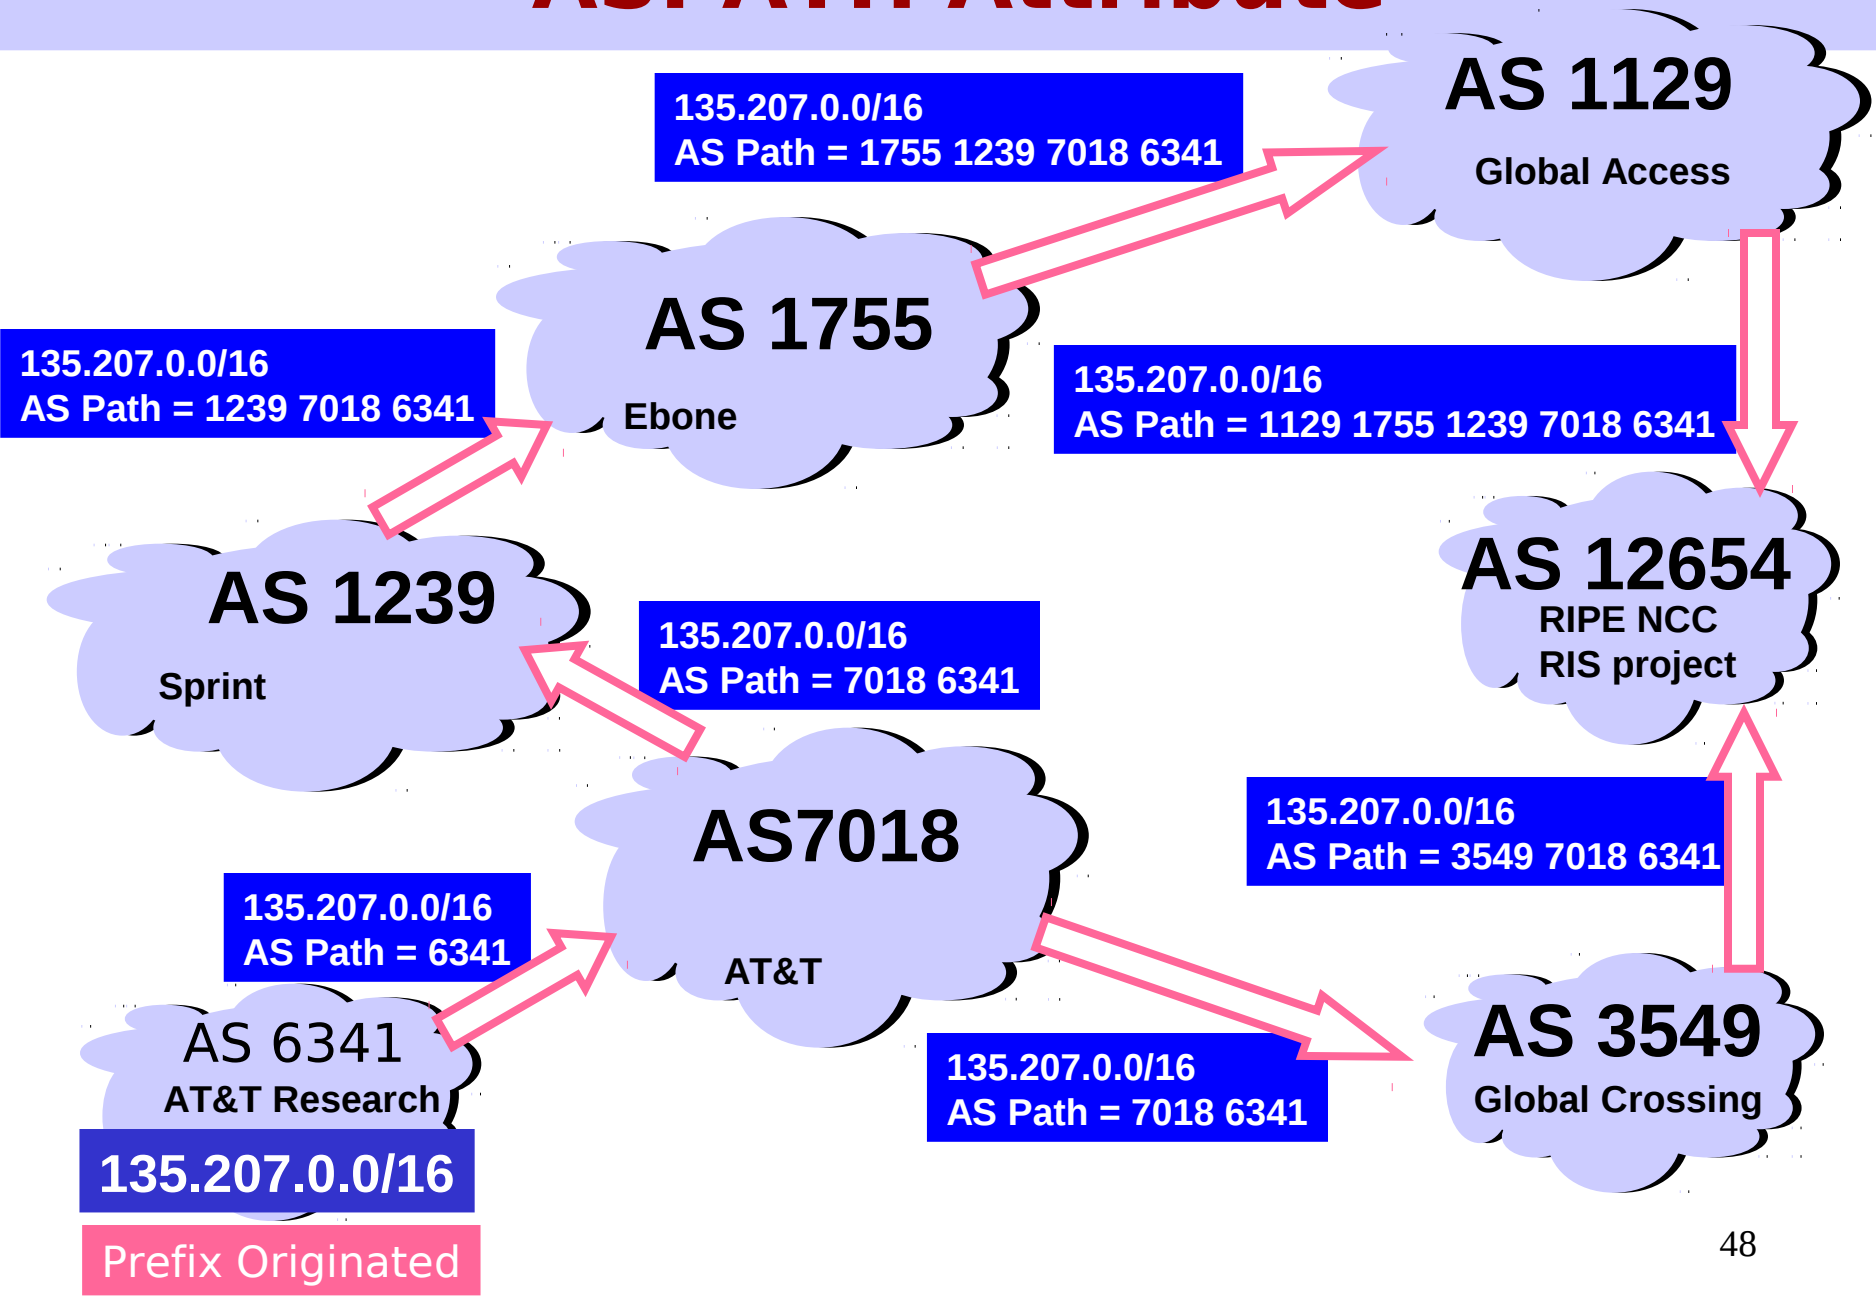

Attribute: Multi-Exit Discriminator (MED)

- When AS's connected via 2 or more links
- AS announcing prefix sets MED
- Enables AS2 to indicate its preference
- AS receiving prefix uses MED to select link, AS1
- Lower MED is better
- Could be used for load balancing traffic

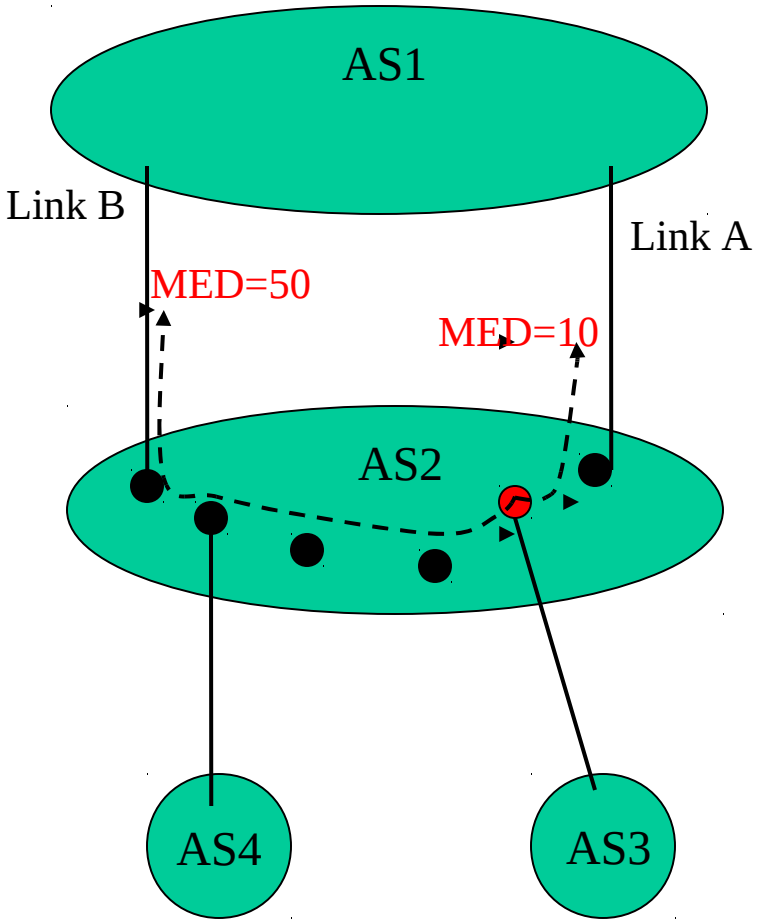

Local Preference

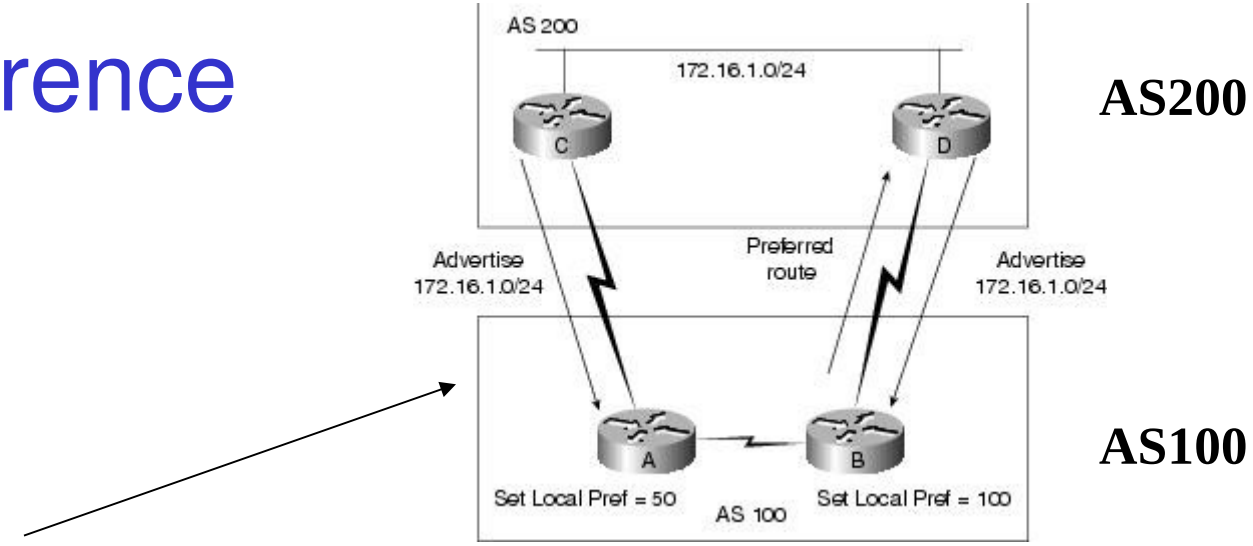

- **AS 100** is receiving two advertisements for network 172.16.1.0 from **AS 200**
  - Router A receives advertisement for network 172.16.1.0,
    - Local preference is set to 50
  - Router B receives advertisement for network 172.16.1.0,
    - Local preference is set to 100
  - Because Router B has a higher local preference than Router A, Router B will be used as exit point from AS 100 to reach network 172.16.1.0 in AS 200

BGP Path Selection

- Example of the complexity in BGP
- BGP uses the following criteria, in the order presented, to select a path for a destination:
  - If the path specifies a next hop that is inaccessible, drop the update.
  - Prefer the path with the largest weight.
  - If the weights are the same, prefer the path with the largest local preference.
  - If the local preferences are the same, prefer the path that was originated by BGP running on this router.
  - If no route was originated, prefer the route that has the shortest AS_path.
  - If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).
  - if the origin codes are the same, prefer the path with the lowest MED attribute.
  - If the paths have the same MED, prefer the external path over the internal path.
  - If the paths are still the same, prefer the path through the closest IGP neighbor.
  - Prefer the path with the lowest IP address, as specified by the BGP router ID.
- Kind of complicated as you can see ….
- But, very flexible, a lot of choices for AS administrators

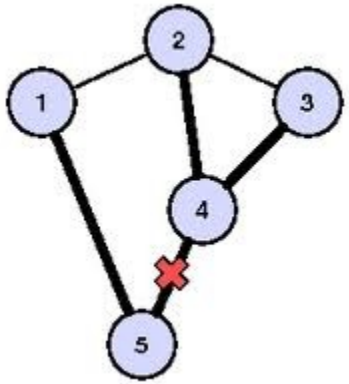

BGP Route Failures

BGP Session Failure

- BGP runs over TCP
  - BGP only sends updates when changes occur
  - TCP doesn't detect lost connectivity on its own
- Detecting a failure
  - Keep-alive: 60 seconds
- Reacting to a failure
  - Discard all routes learned from the neighbor
  - Send new updates for any routes that change

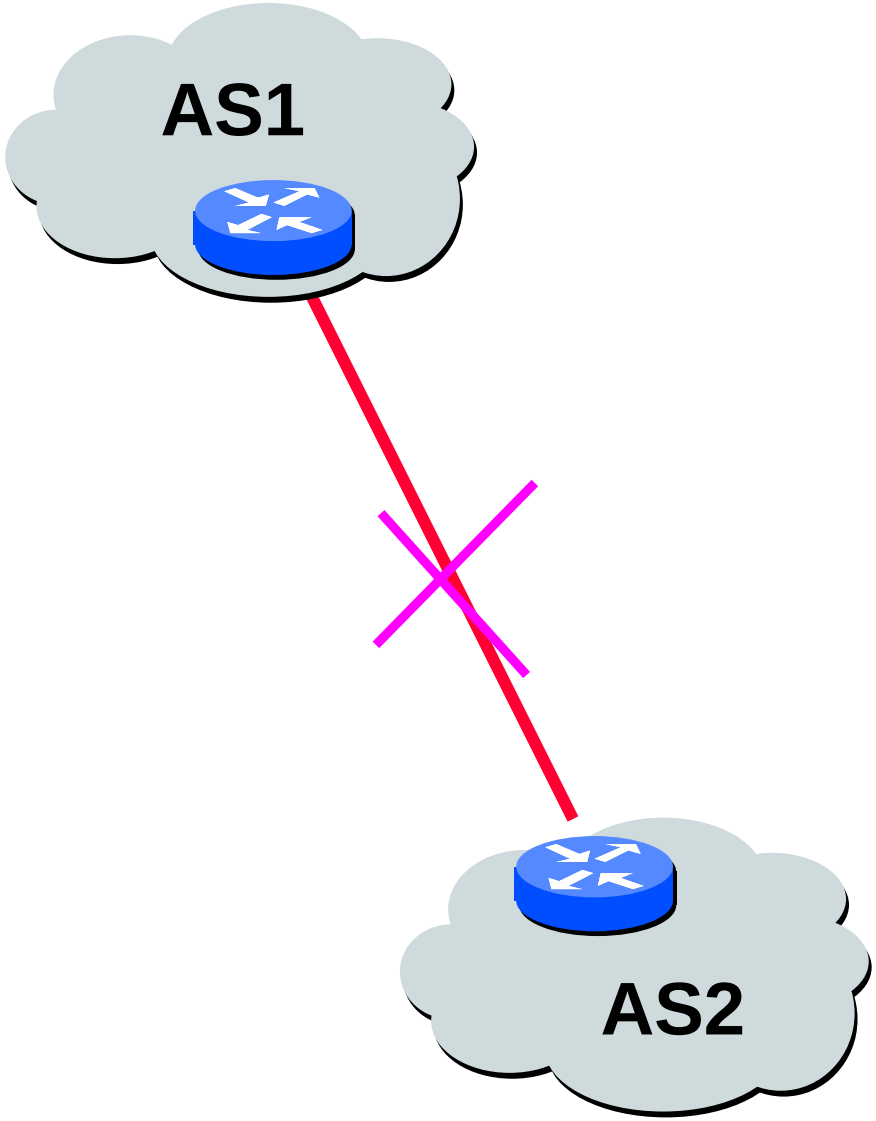

Routing Change: Before and After

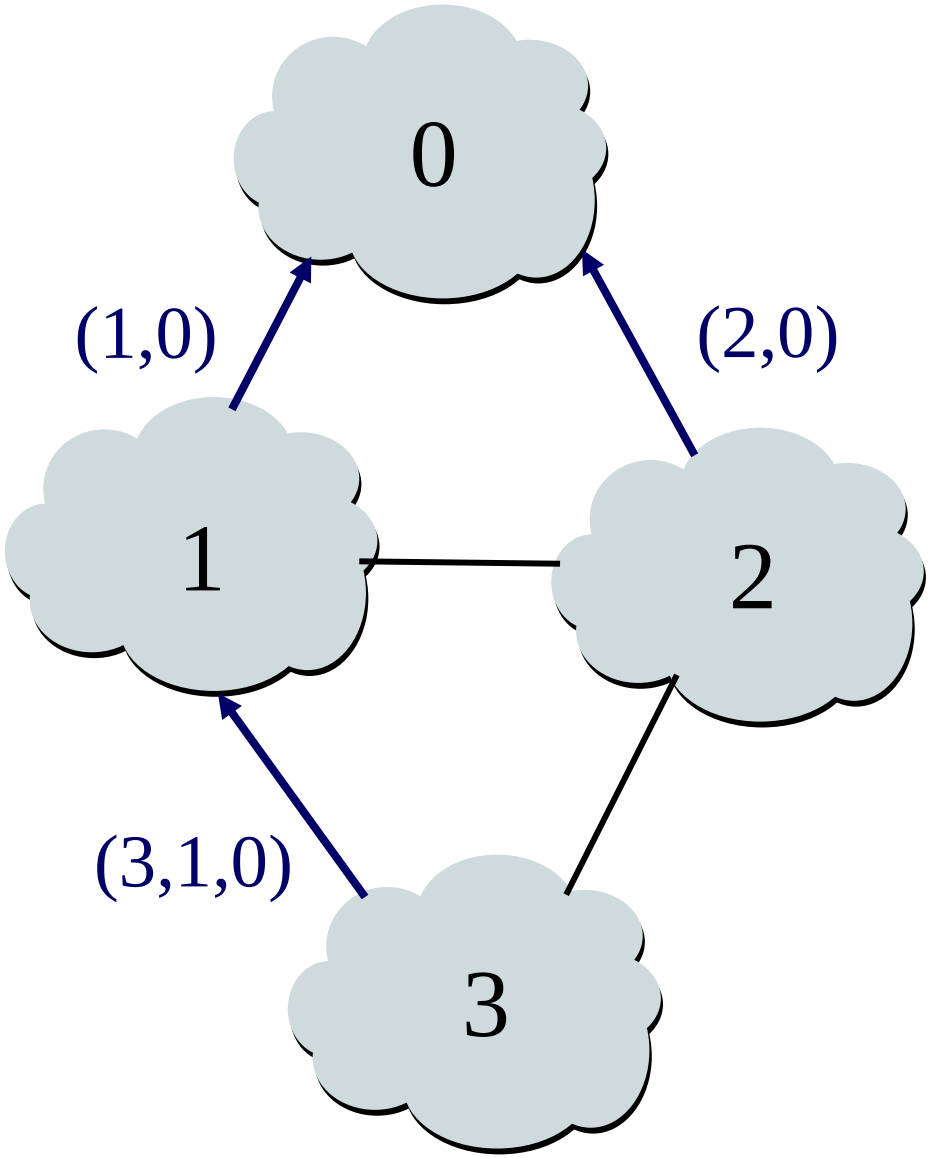

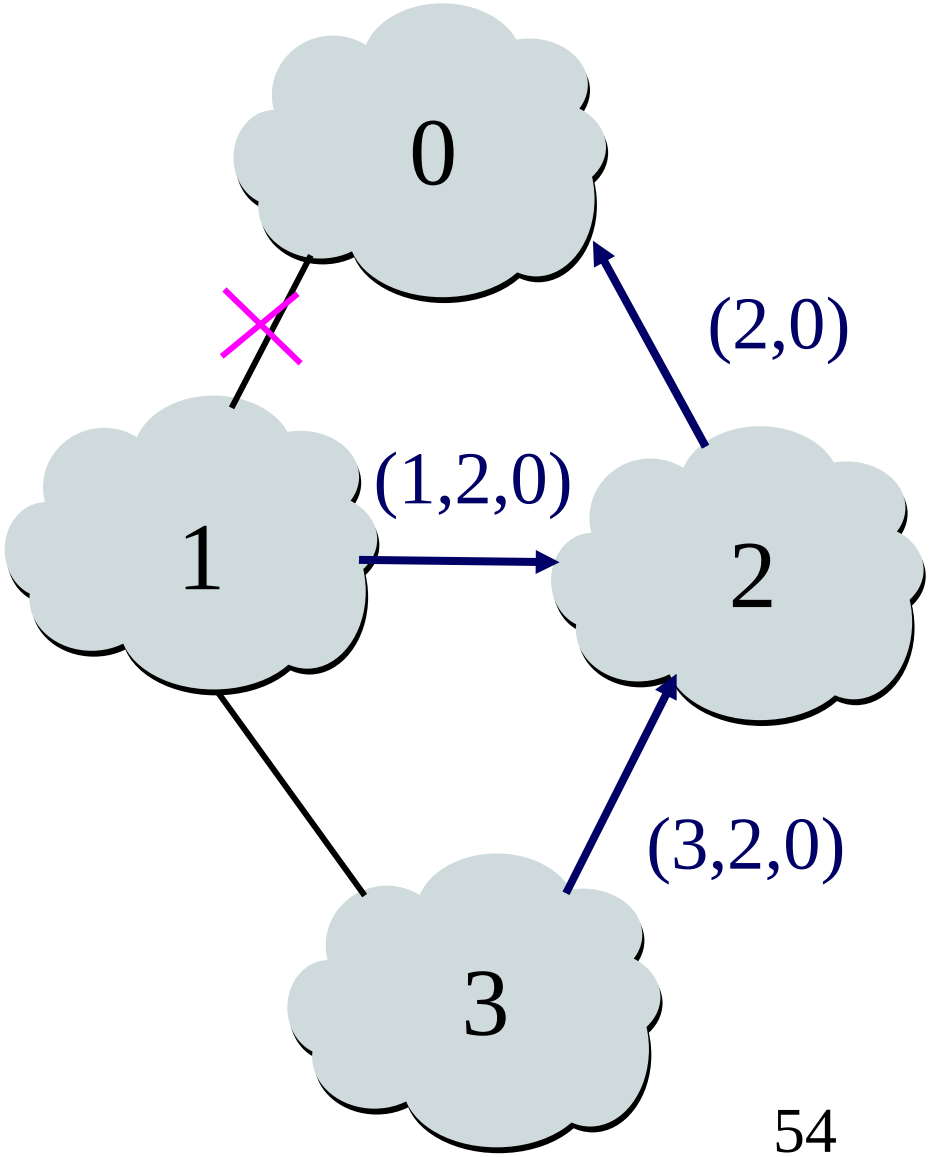

Routing Change: Path Exploration

- AS 1
  - Delete the route (1,0)
  - Switch to next route (1,2,0)
  - Send route (1,2,0) to AS 3
- AS 3
  - Sees (1,2,0) replace (1,0)
  - Compares to route (2,0)
  - Switches to using AS 2

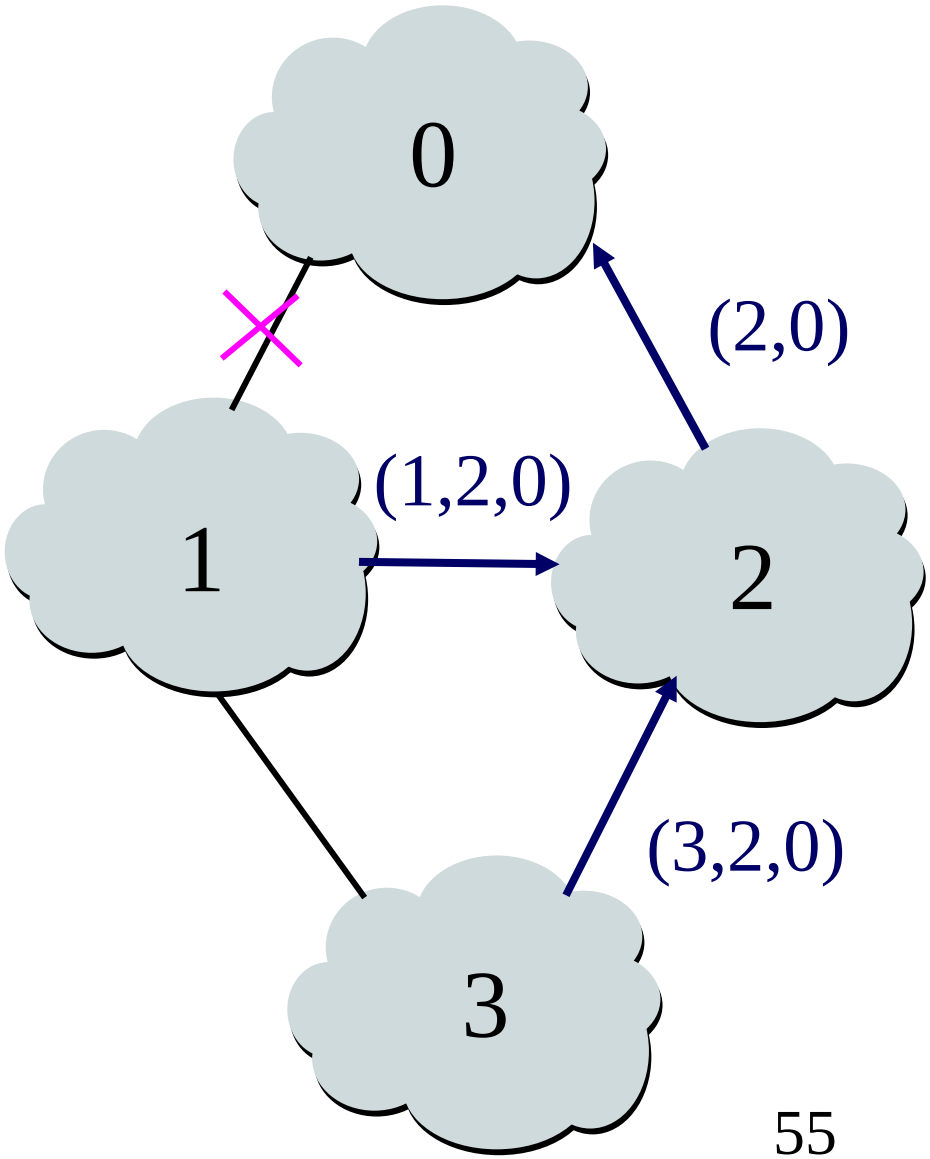

BGP Converges Slowly

- Path vector avoids count-to-infinity
  - But, ASes still must explore many alternate paths
  - … to find the highest-ranked path that is still available
- Fortunately, in practice
  - Most popular destinations have very stable BGP routes
  - And most instability lies in a few unpopular destinations
- Still, lower BGP convergence delay is a goal
  - Can be tens of seconds to tens of minutes
  - High for important interactive applications
  - … or even conventional application, like Web browsing

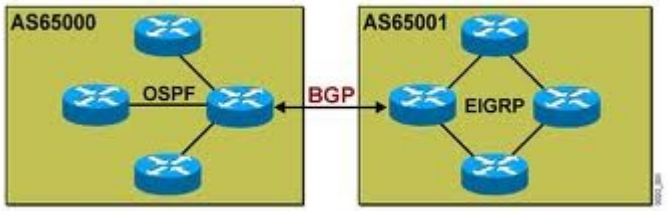

BGP Problems

Route Flapping

- BGP peers exchange routes, send updates
- When route is repeatedly advertised and withdrawn,
  - Said to be **'flapping'**
- Flapping routes cause instability in Internet routing table
- Cisco routers running BGP contain optional mechanism designed to dampen destabilizing effect of flapping routes

- What does it mean to dampen a route?

Route Flapping

- BGP process assigns a penalty of **1000** each time it flaps
- When penalty value exceeds a set limit,
  - Route is moved into 'historical' list of routes, dampened, and suppressed for 15 minutes
- Maximum suppress limit is, one hour

BGP Security Problems

BGP Security Problems

- Infrastructure built on top of BGP is highly vulnerable to malicious attacks
- Fundamentally, no way to guarantee BGP router uses allocated AS number or that it holds address space it advertises
- Most of BGP's security problems are result of lack of verification that IP address space belongs to a given AS

Border Gateway Protocol

- Attacker Goals
  - Why attack BGP?
  - What are advantages?

Border Gateway Protocol

- Attacker Goals
  - Why attack BGP? What advantages?
  - Black Hole
    - Drop traffic, make a prefix unreachable
    - Attract traffic to a router then drop it

  - Redirection
    - Traffic flowing to a particular network forced to take different path, may cause link to collapse

BGP

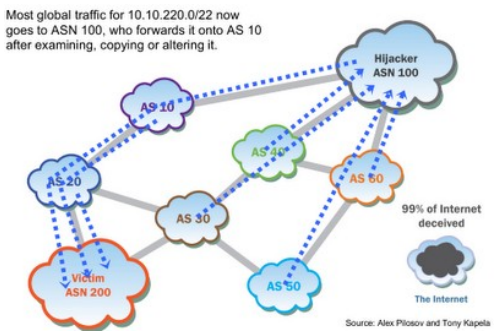

- Why attack continued
  - Eavesdrop or Modify
    - Pass data through link to eavesdrop or modify data
  - Instability
    - Cause route dampening, connection outages
      - Routes that change too frequently get penalized

        "Route Flapping" leads to

        Route dampening - routes assigned a less preferred status
    - Cause increased BGP traffic and cause route convergence delays

BGP

- How to attack BGP?
  - Provide wrong information
    - Connections that don't really exist
    - Reroute traffic through compromised routes
    - Provide contradictory or confusing information
  - Provide more frequent information
    - Advertise routes more often
    - Destabilize routing tables
  - Example follows ...

YouTube Gets DoS'd

- Feb. 2008, Pakistan government bans YouTube - blasphemous content
- Nobody from Pakistan can get to YouTube
- PCCW,
  - One of the largest communications providers for Pakistan and China, was supposed to only block Pakistani users … blocked all users from YouTube,
  - Not just the Pakistani ones ...

YouTube Gets DoS'd

- Result ... all BGP speaking routers believed Pakistan Telecom provided best connectivity to YouTube

- A complete denial of service (DoS),
  - Intentional or not!!!

BGP Routing Details

- BGP rules state that longer routes are more specific and preferred, more bits for network portion
- So, YouTube, owns IP space
  - 208.65.153.0/24,
  - 208.65.152.0/24 and
  - 208.65.154.0/23,
- YouTube announces single aggregated BGP route for /24 prefixes, announced as 208.65.152.0/22

208.65.152.0/22 via AS 36561 (YouTube)

208.65.153.0/24 via AS 17557 (Pakistan Telecom)

Hijacked YouTube Visuals

- RIPE NCC has tools that monitor BGP routes
  - RIPE is regional Internet registry for Europe, Middle East and Central Asia

Have an outdated Page That Summarizes the Attack

http://www.ripe.net/news/study-youtube-hijacking.html

On Youtube

https://www.youtube.com/watch?v=IzLPKuAOe50#t=19

- Actual animation of entire event complete with music!!!!

More BGP Problems

Blackhat 2008
Tony" Kapela and
Alex Pilosov

Similar BGP problem revealed at Blackhat 2008

- Anyone with BGP router ...
  - ISPs, large corporations, governments,
  - Could intercept data headed to a target IP address or group of addresses
- Attack intercepts only traffic headed to target addresses, not from them

2008 Demo at Blackhat

- Tony Kapela and Alex Pilosov
- Man-in-the-middle attack demonstrated at

Defcon 2008

- Redirected traffic bound for Defcon to system they controlled in New York and then routed it back to Las Vegas
- Good analysis of this attack below
- While BGP eavesdropping has long been a known weakness but no one was known to have intentionally exploited it until this proof of concept

http://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf

2008 Demo at Blackhat

- What did they discover about BGP?
- Pilosov's innovation
  - Forward intercepted data to actual destination, so that no outage occurs
  - Use these AS's to forward stolen data to its rightful recipients
  - Used protocol against itself to subvert it !!!

2008 Demo at Blackhat

- What could you do with this attack?
- Corporate espionage,
- Nation-state spying or
- Intelligence agencies looking to mine Internet data
- Don't need cooperation of ISP's …

Good Overview of BGP Hijacking

https://www.bishopfox.com/blog/2015/08/an-overview-of-bgp-hijacking/

BGP Vulnerabilities

Recap ... vulnerabilities that allow these types of attacks to happen

- Lack of authentication of BGP updates
  - Are they coming from “trusted” routers?
- Updates sent in the clear
- Updates themselves can be bogus
  - By accident or deliberate can poison the routing tables

Memo on BGP Security Vulnerabilities Analysis

http://www.ietf.org/rfc/rfc4272.txt

Current BGP Attacks

- **2013 Renesys, Provides Internet Intelligence**
- February 2013, observed sequence of events, lasting from few minutes to several hours in duration, in which global traffic was redirected to Belarusian ISP GlobalOneBel
    - These redirections took place on an almost daily basis throughout February
- Affected countries included the US, South Korea, Germany, the Czech Republic, Lithuania, Libya, and Iran
- Another set of traffic hijack incidents took traffic to Iceland

http://www.renesys.com/2013/11/mitm-internet-hijacking/

BGP Fixes

- Countermeasures
  - TCP connection hijack protection
    - MD5 hash signature
    - Insure that BGP messages have source address of legitimate peering BGP speaker
    - Absolutely, identify BGP router

BGP Fixes

- Route Filtering
  - Used to enforce business relationships between AS's
  - Create Access Control Lists (ACL's) of prefixes for sending/receiving updates
  - Egress filters allow control of announced routes to peers
  - Ingress filters check incoming routes for validity
    - Make sure origin AS of route owns prefix

BGP Fixes

- Route Filtering continued
  - What's the Problem ?
    - Hard to keep Internet routing registries current
    - ISP's trust that their peer routers sending correct information
    - Also, in practice filtering is against dynamic nature of Internet
    - Policies change often, structure of AS's not tree
    - AS's have multiple connections, difficult to apply strict filters

One Fix, SBGP

- SBGP – Secure BGP
  - Extension to BGP
  - Protect BGP from malicious or mistaken updates
  - Adds authorization and authentication
    - Attribute added to BGP updates to ensure updates valid
    - Route messages secured with IPSec
  - Based on PKI cryptography

    http://www.net-tech.bbn.com/sbgp/sbgp-index.html

SBGP

- SBGP – Secure BGP
  - Adds Address Attestation (AA)
    - Verify origin AS is authorized to advertise a particular address block
    - Verify AS owns that address block
  - Adds Route Attestation (RA)
    - Authorize neighbor AS's to propagate route contained in an update

SBGP

- SBGP
  - More details
    - Uses PKI to authorize AA's and RA's
    - Private keys stored in S-BGP speakers
    - Public keys made available by hierarchical PKI infrastructure

    - Any problems with this?

Problems with SBGP

- Need to have hierarchical PKI in place and trusted by **all** ISP's
- Crytography intensive and part of huge overhead when BGP router reboots
- Routers may need large memory 20 MB to store public keys
- Routers can't always sign routes if routes have been aggregated
  - Routes will have come from multiple sources

Problems with SBGP

- Have prevented SBGP from being deployed
- Alternative methods have been suggested
  - **CISCO proposed soBGP** – Secure Origin BGP
  - Lightweight alternative to SBGP
  - Uses existing trust relationships to validate certificates - "Web of Trust"
  - **IRV** – Companion protocol to BGP
    - Uses IRV servers,
    - Updates are verified by each AS IRV server in AS-PATH

Router Table Growth

- BGP routing tables are master lists of network destinations stored in backbone routers
  - Used to determine best available path between networks
- Experts worried about explosive growth in BGP routing tables
  - Strains processing and memory requirements of Internet's core routers
  - BGP table growth drives up carrier costs
  - Everyone worries about costs!

Routing Table Sizes

- As of 2009, BGP routing table had around 240,000 routes
- Trend over last year and a half has been exponential growth
- Experts worried that some older routers could fail

BGP Router Table Entries by Year to 2007

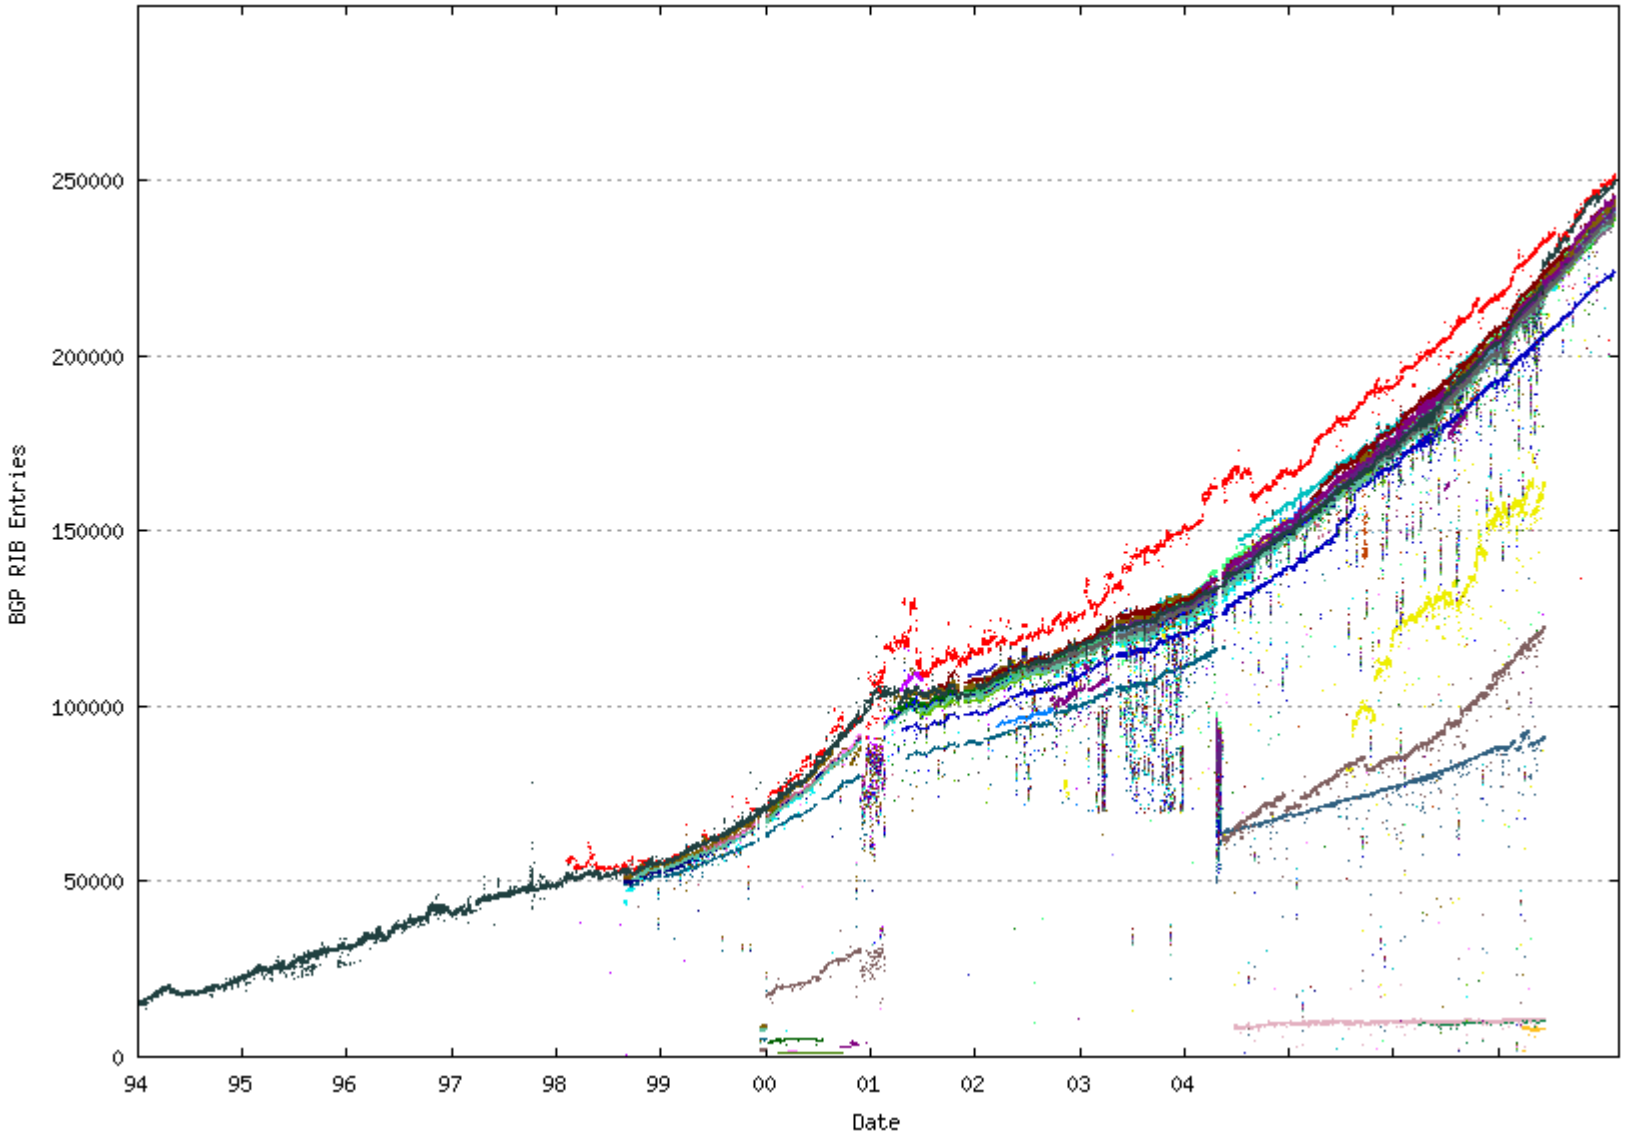

BGP Router Table Entries by Year to 2011

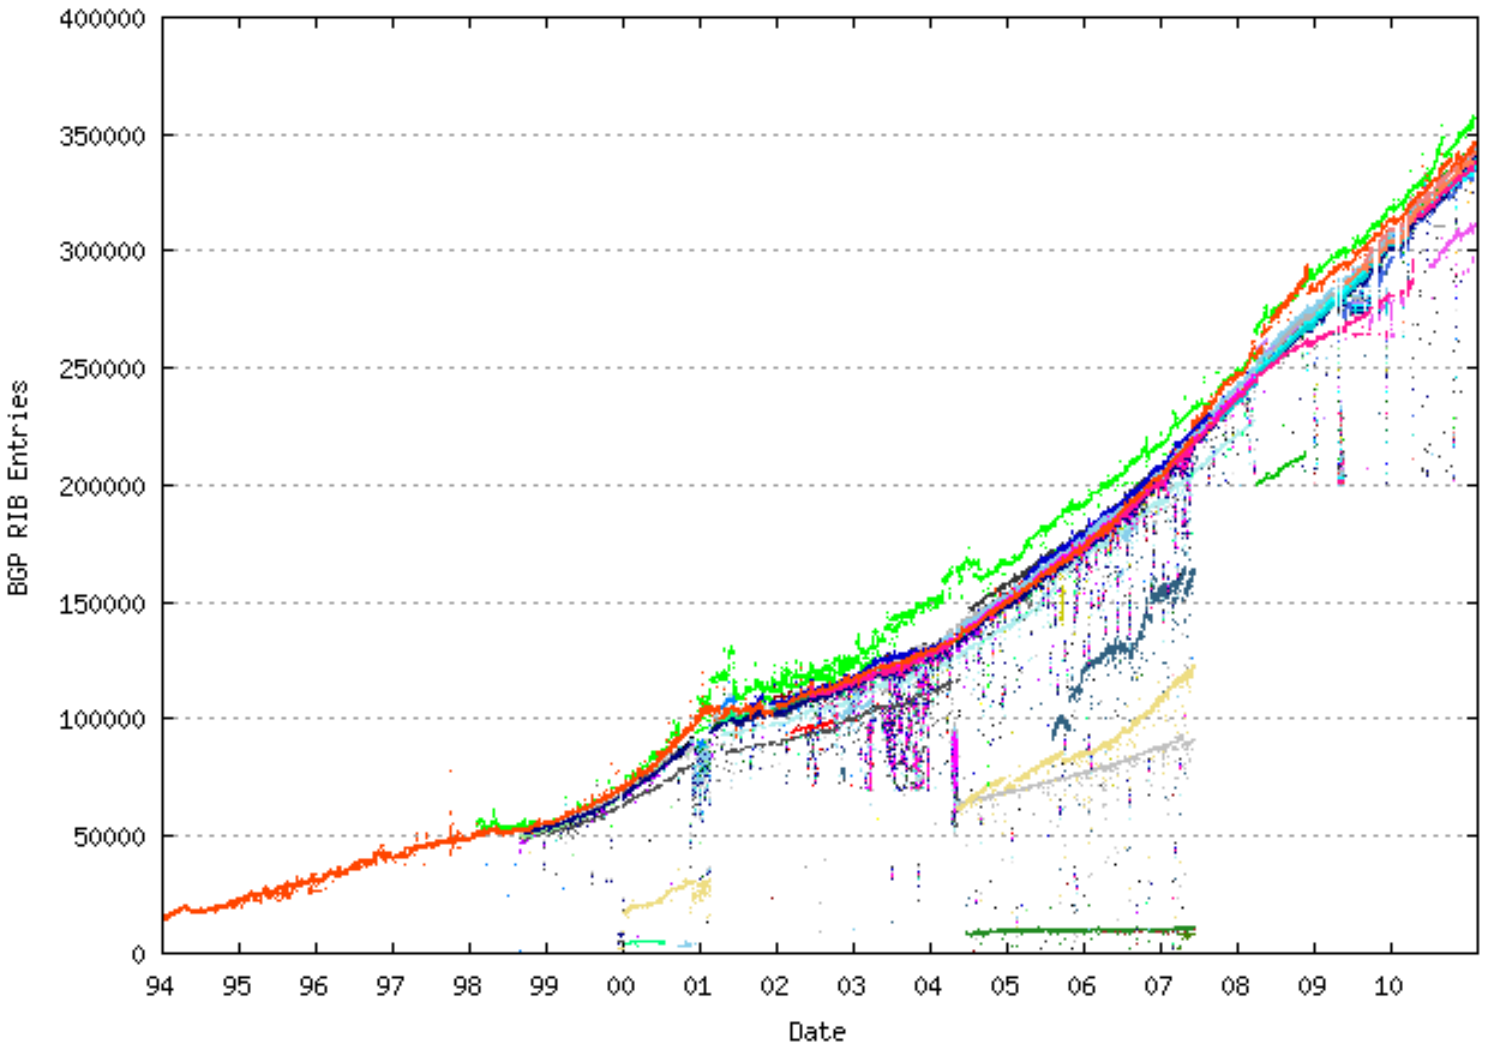

Solutions from Internet Research Task Force (IRTF)

- Recommend keeping BGP, tweaking it so carries different, preferably less, information

- Any solution that the Routing Research Group comes up with is at least five years away from commercial availability, experts agree.
  - IRTF cautiously optimistic about group's ability to solve the routing table growth problem

- Reference:

  http://www.infoworld.com/archives/emailPrint.jsp?R=printThis&A=/article/07/09/27/radical-rethink-of-internet-routing_2.html

Summary

- Hierarchy continues as a design principle at the Global Internet
  - Networks organized into Autonomous Systems
  - Independently managed with independent connections to other AS's
- Routing between them using BGP
- Routing becomes not just an exercise in shortest path delivery
  - But, incorporates policy decisions between ISP's based on political (monetary) alliances

References

- Network Peering and Exchange Points

  http://www.infocellar.com/networks/internet/nap-ixp.htm

- Cisco and Networking Planet Overview of BGP

  http://www.ciscopress.com/articles/article.asp?

  p=1565538&seqNum=2

  http://www.enterprisenetworkingplanet.com/netsp/article.php/3615896/Networking-101-Understanding-BGP-Routing.htm

- BGP Statistics

  http://bgp.potaroo.net/

Midterm due Monday, February 27

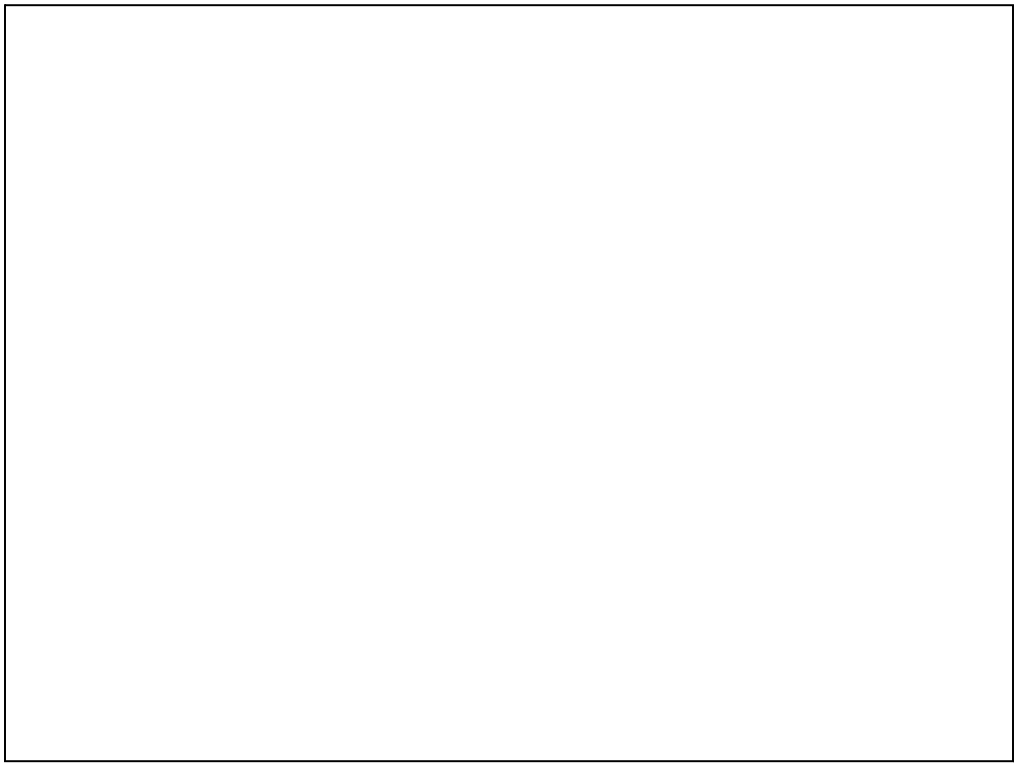

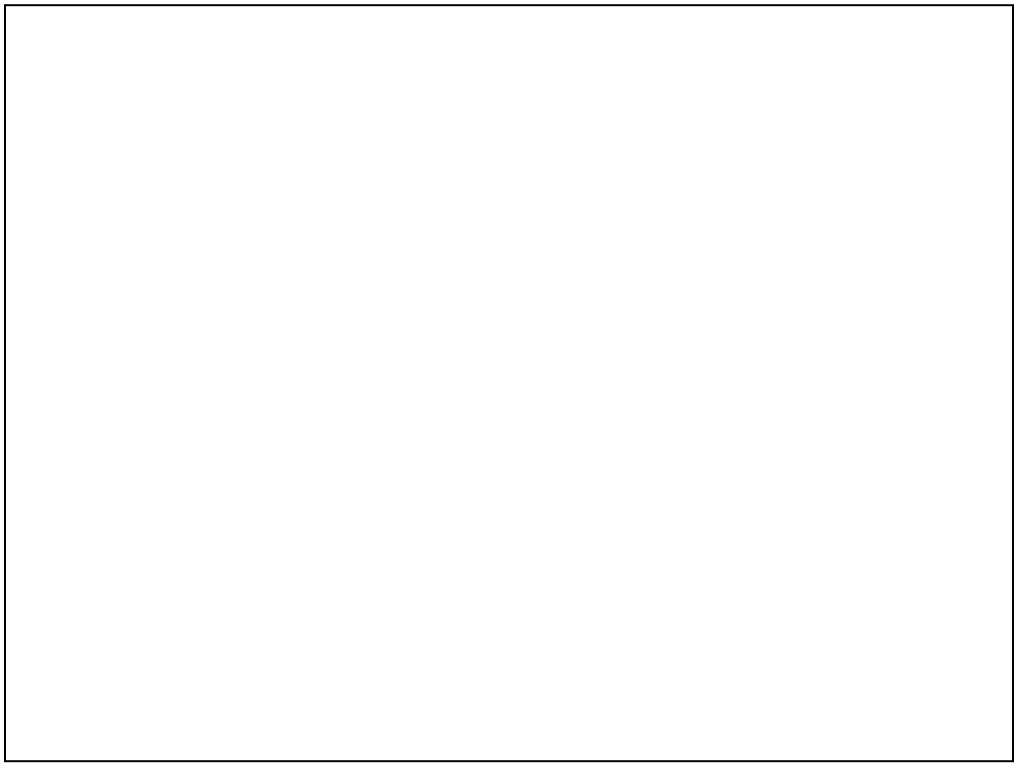

Connected to Internet

- What does it mean to be connected to the Internet?
  - Packets sent to host arrive at host
  - Packets sent back arrive at destination

- Means must have a path to you
- Your ISP must have a path to you
- My IP must lie within an address space that gets advertised as a route by others

- Else no-one can find me

BGP's Role in Connectivity

- Advertises routes
  - Think of these as promises to carry data to IP space represented in route being advertised
- You promise that If someone sends data to route that is advertised
  - Know how to carry data to its ultimate destination
  - Don't want to advertise routes to places that we don't know how to reach !!!

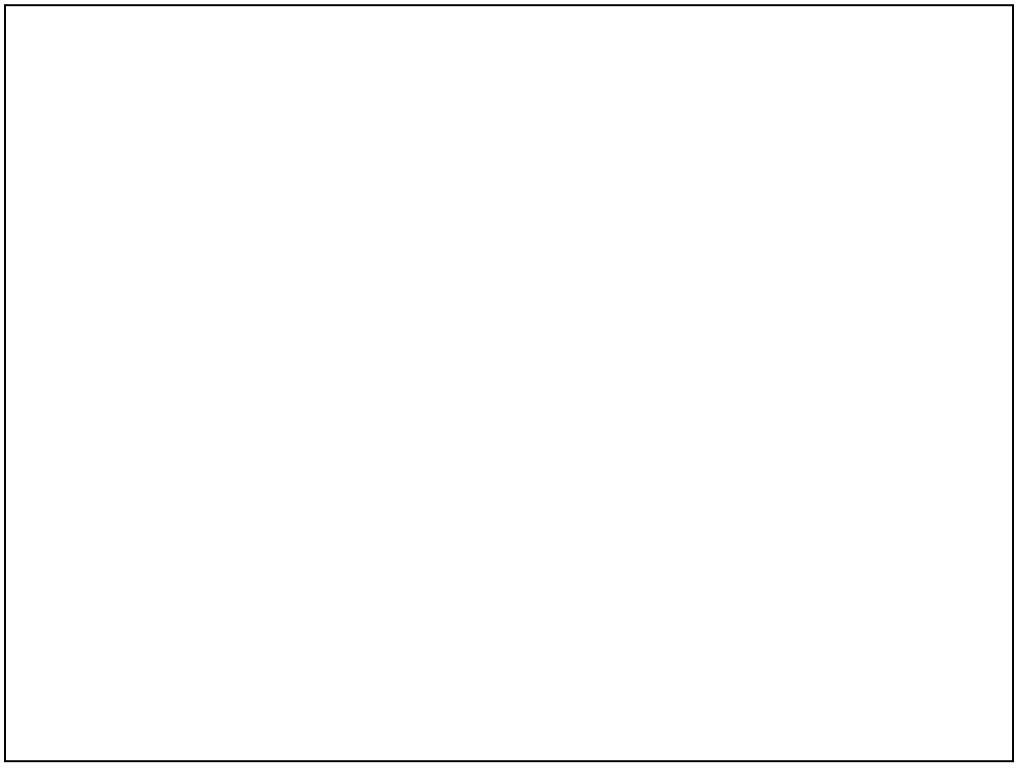

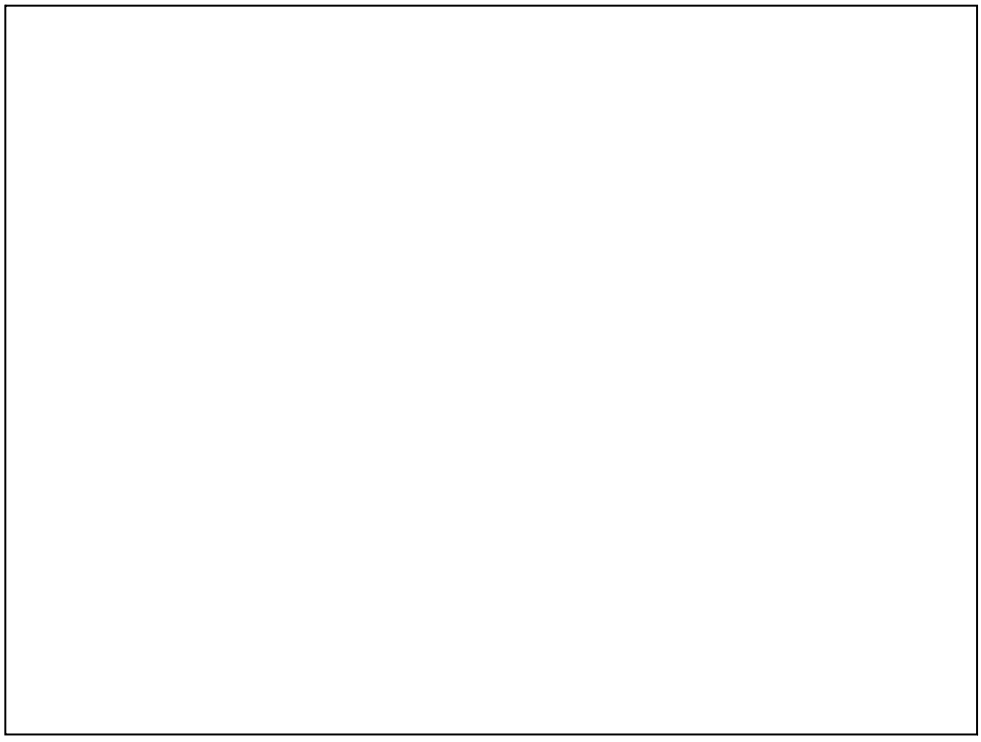

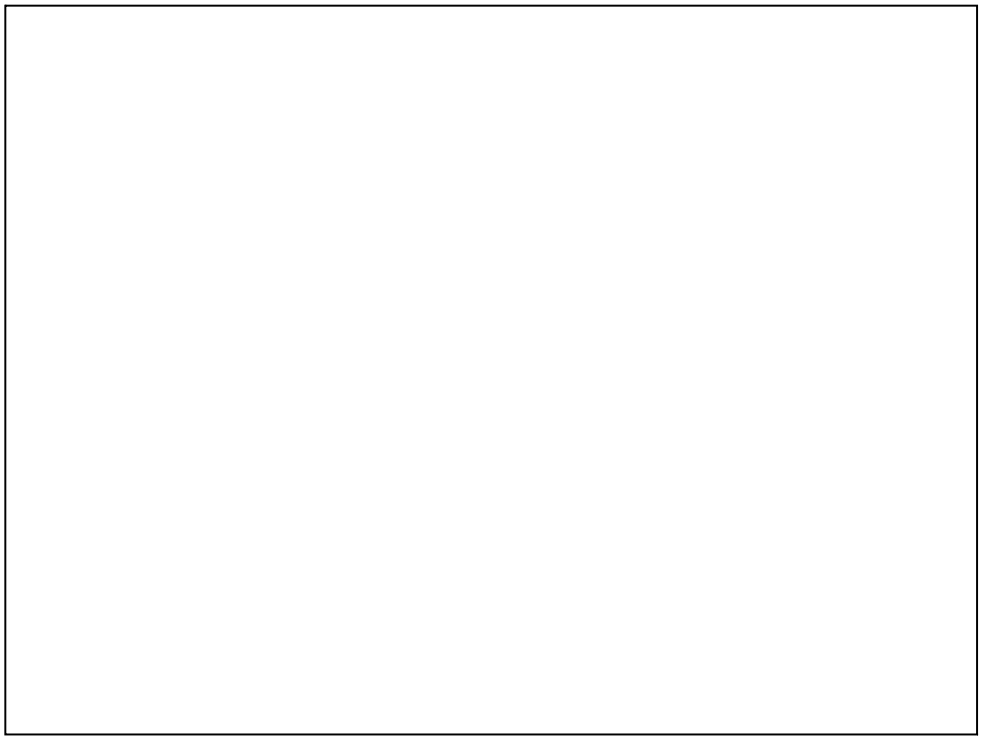

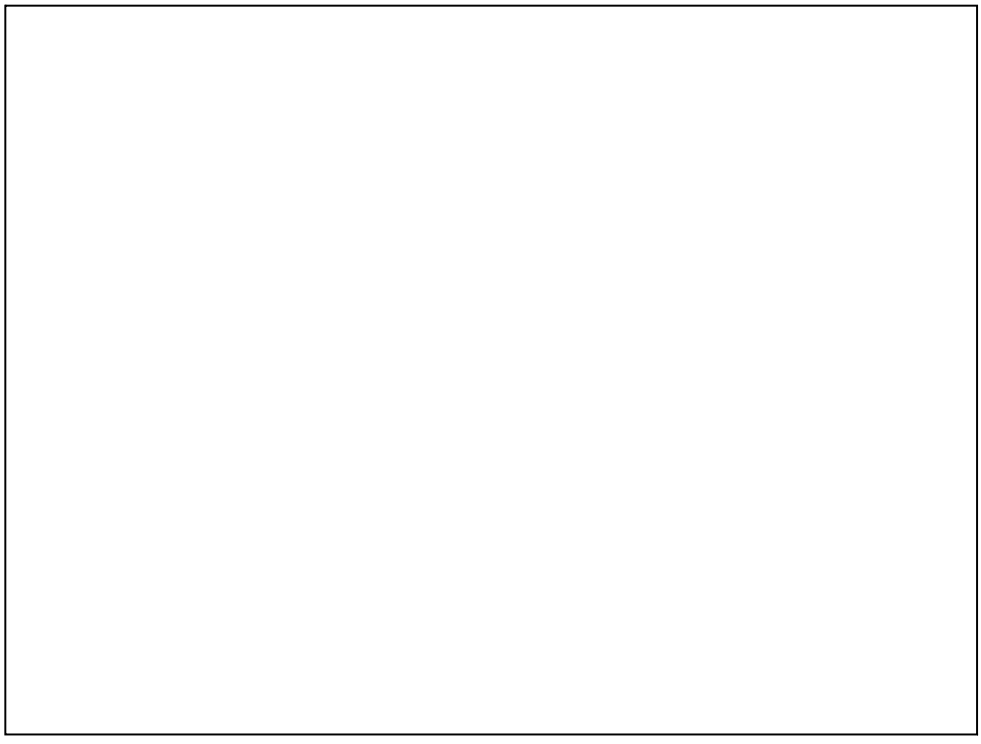

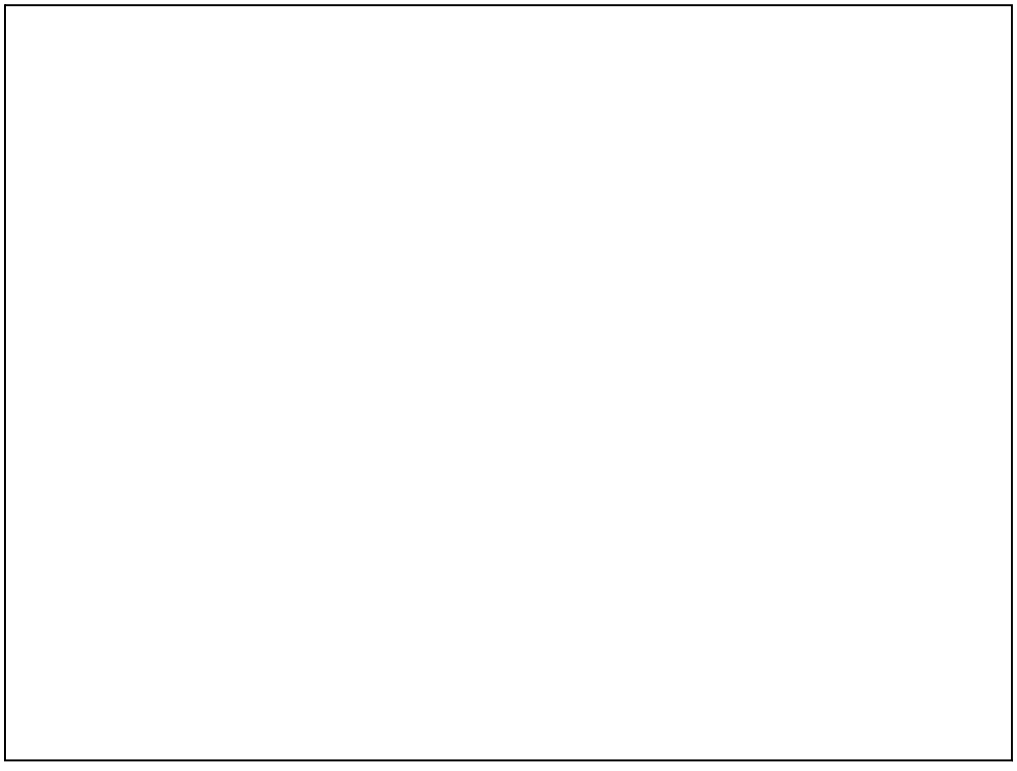

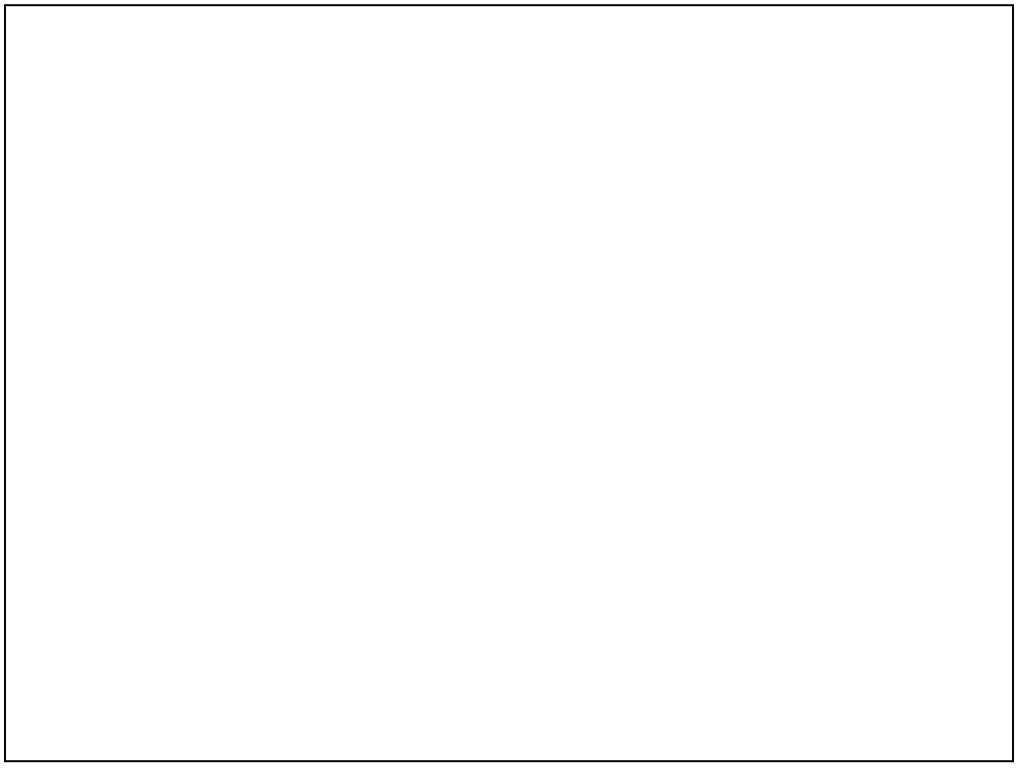

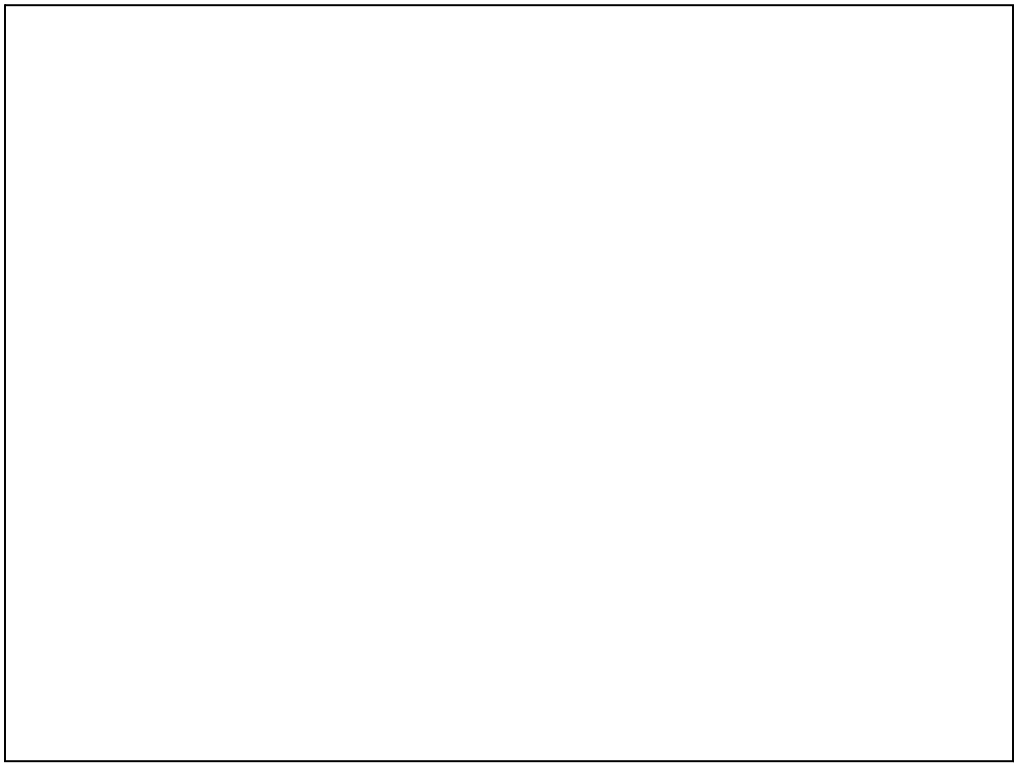

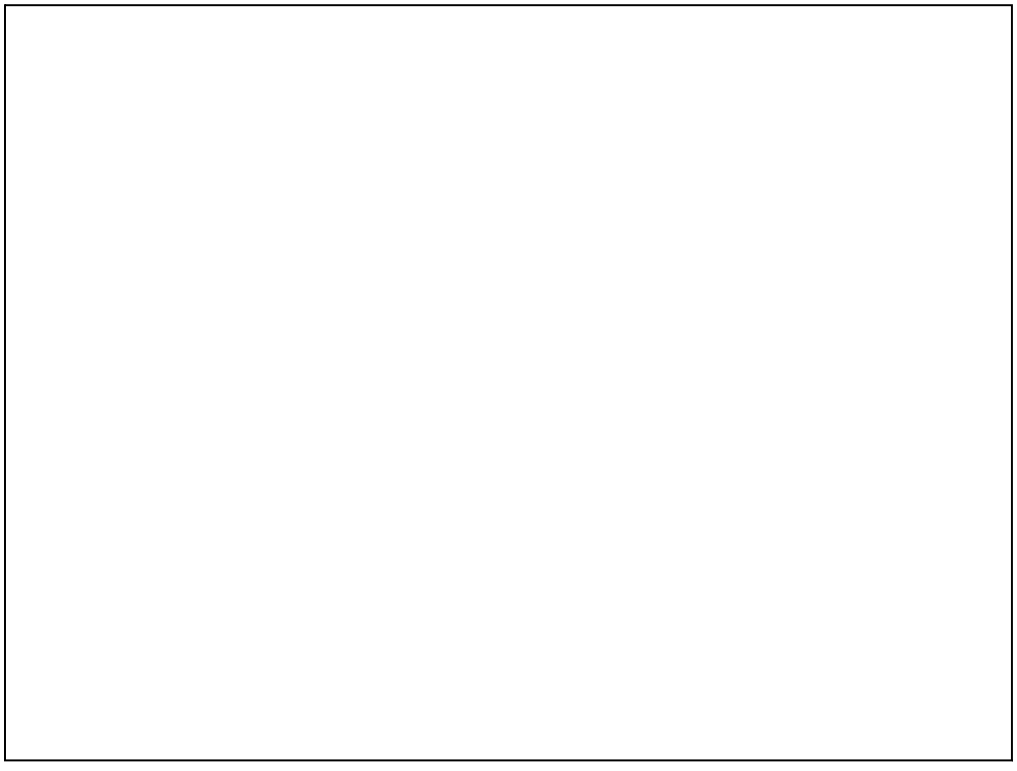

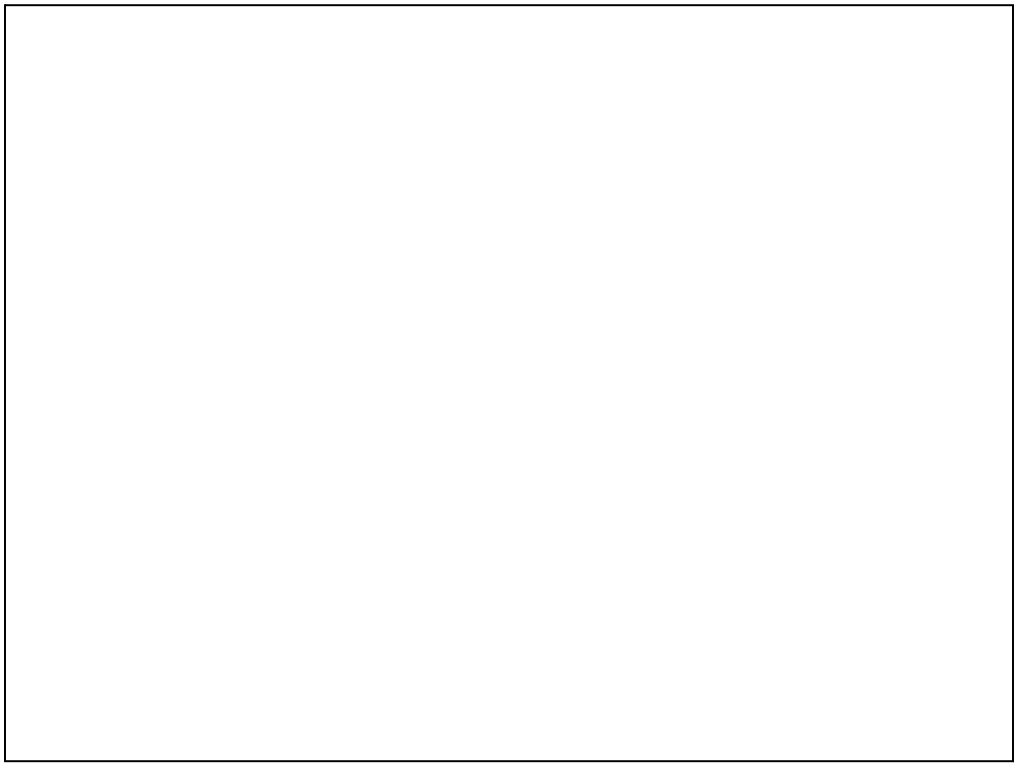

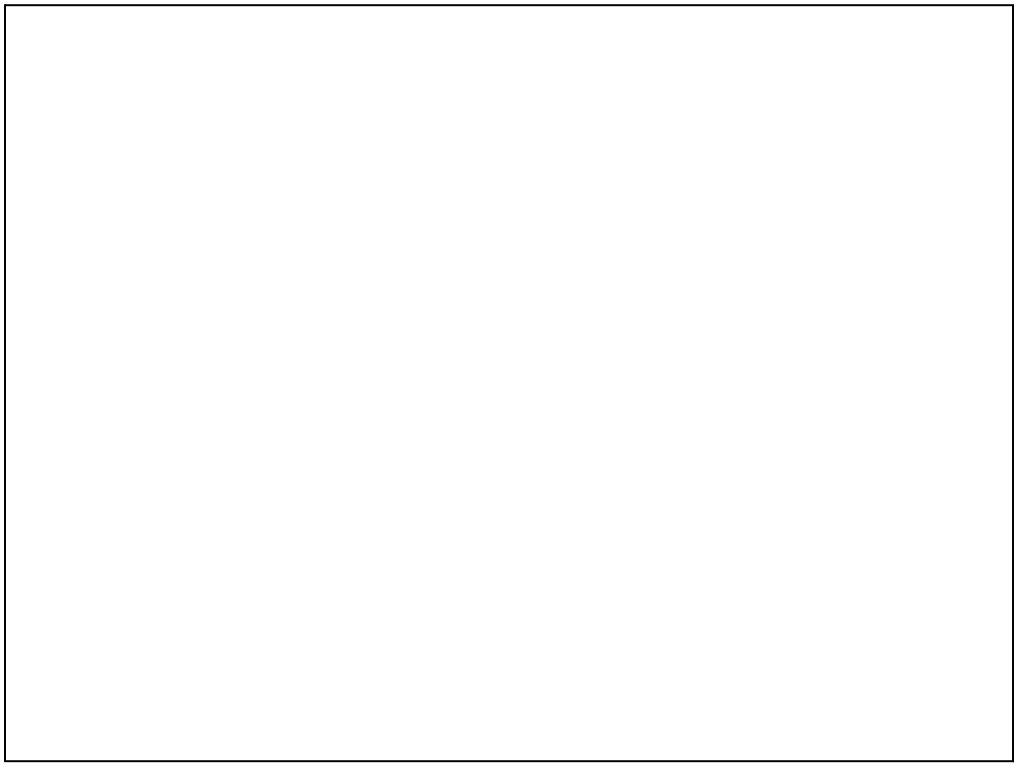

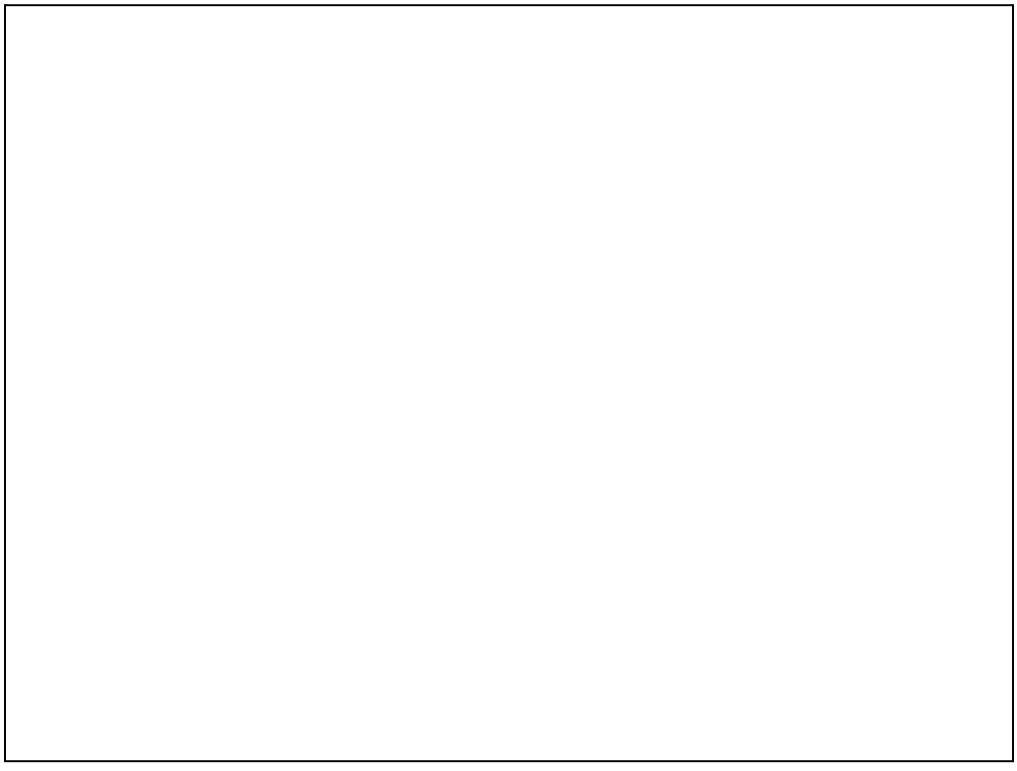

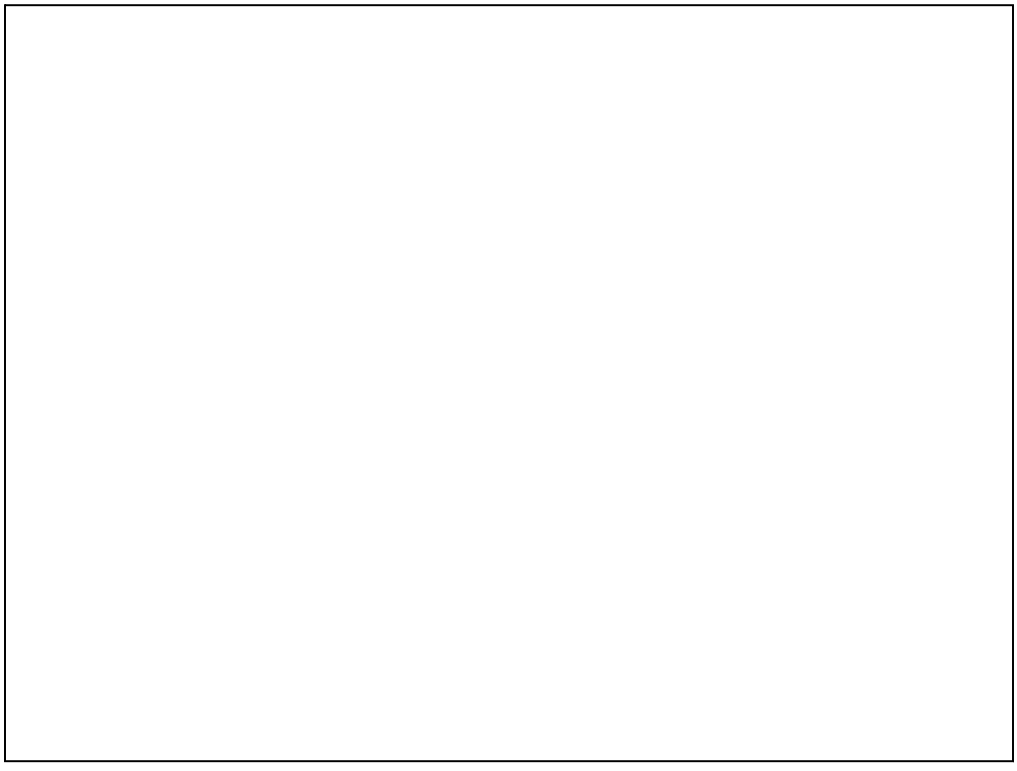

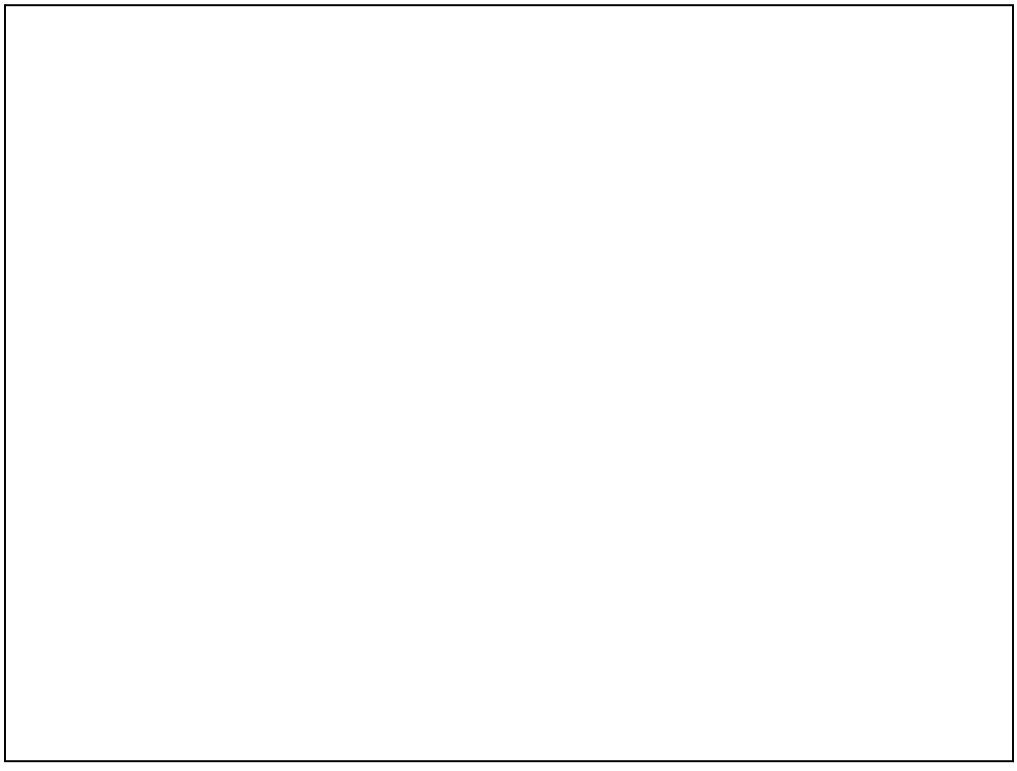

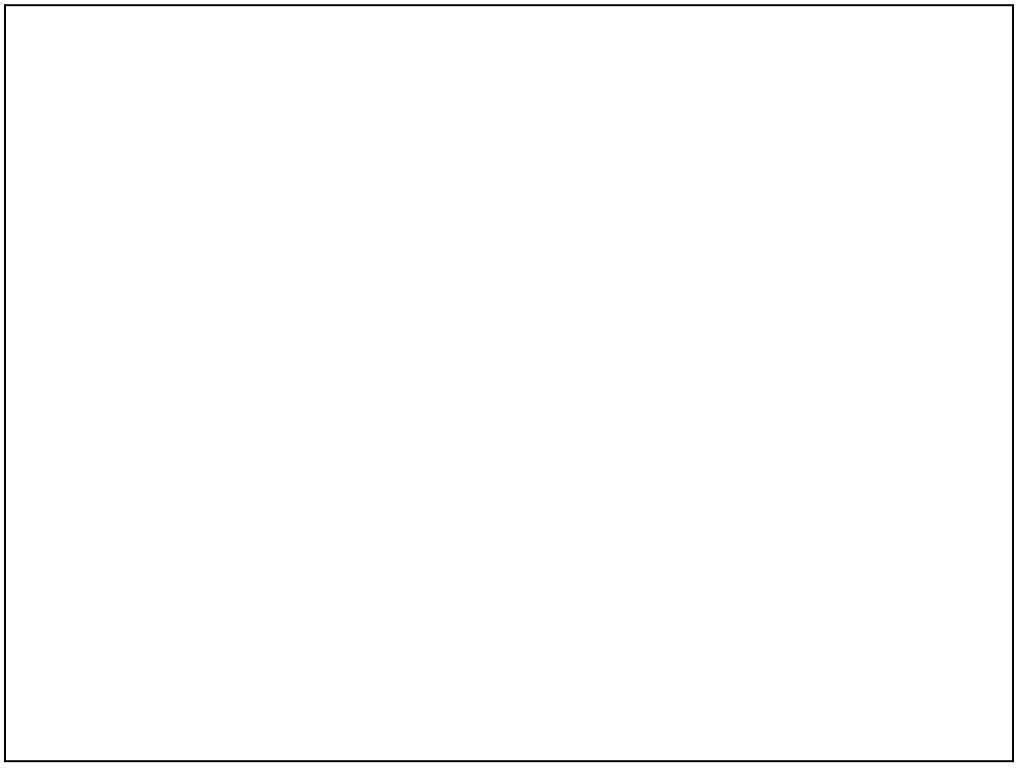

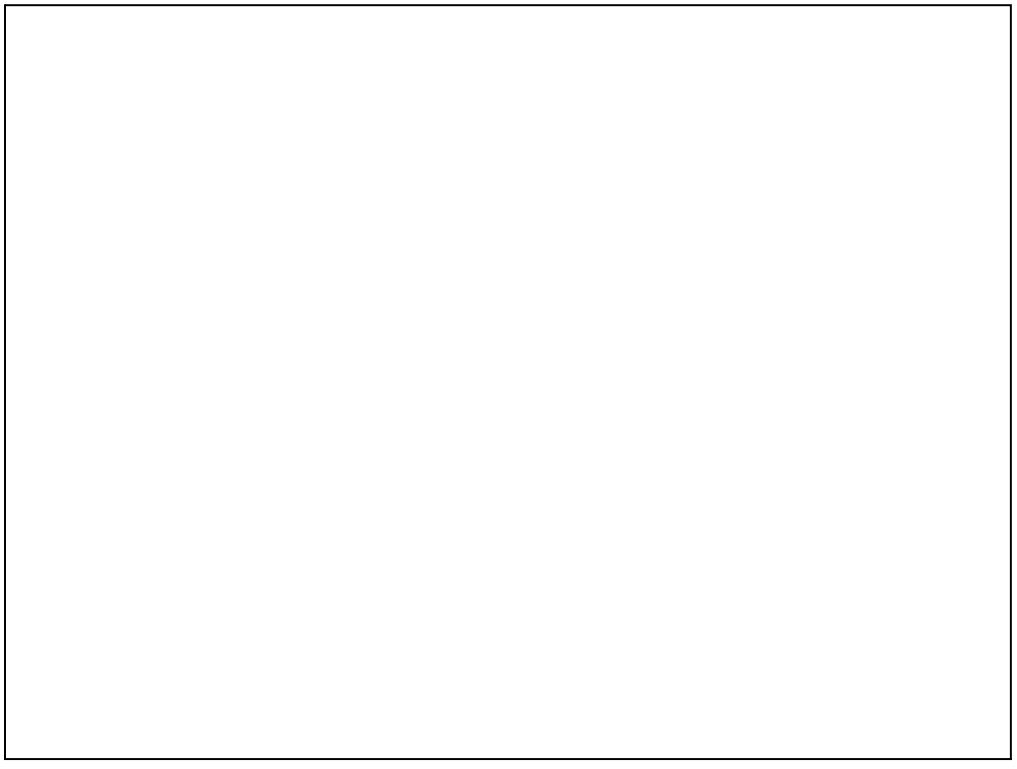

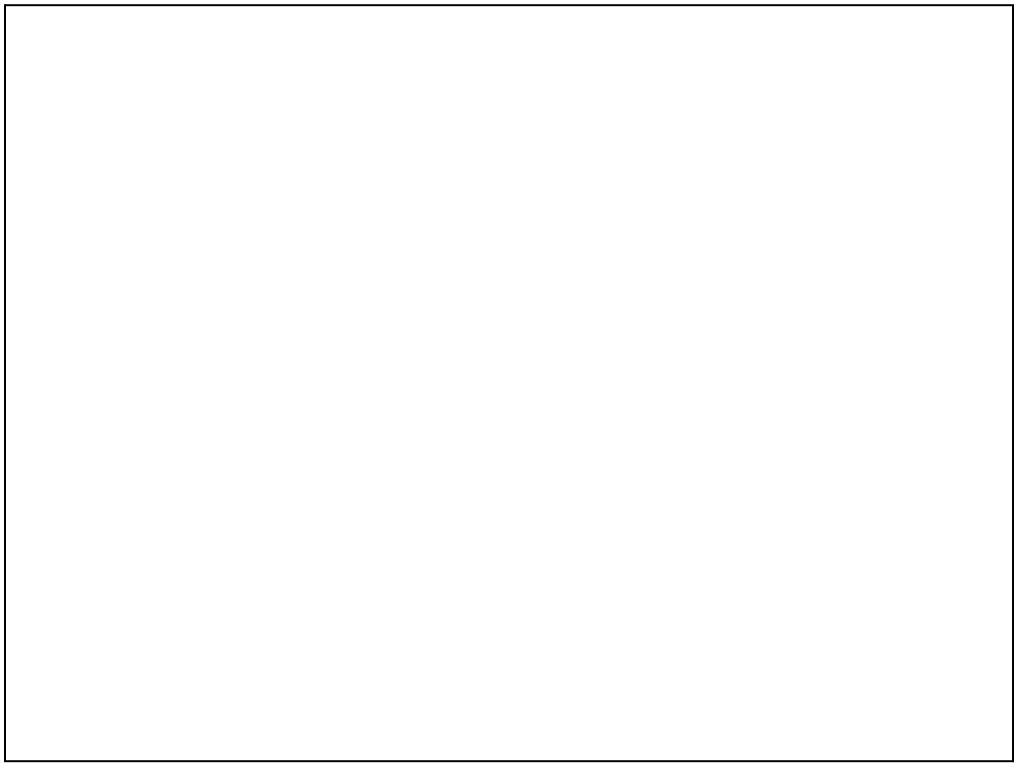

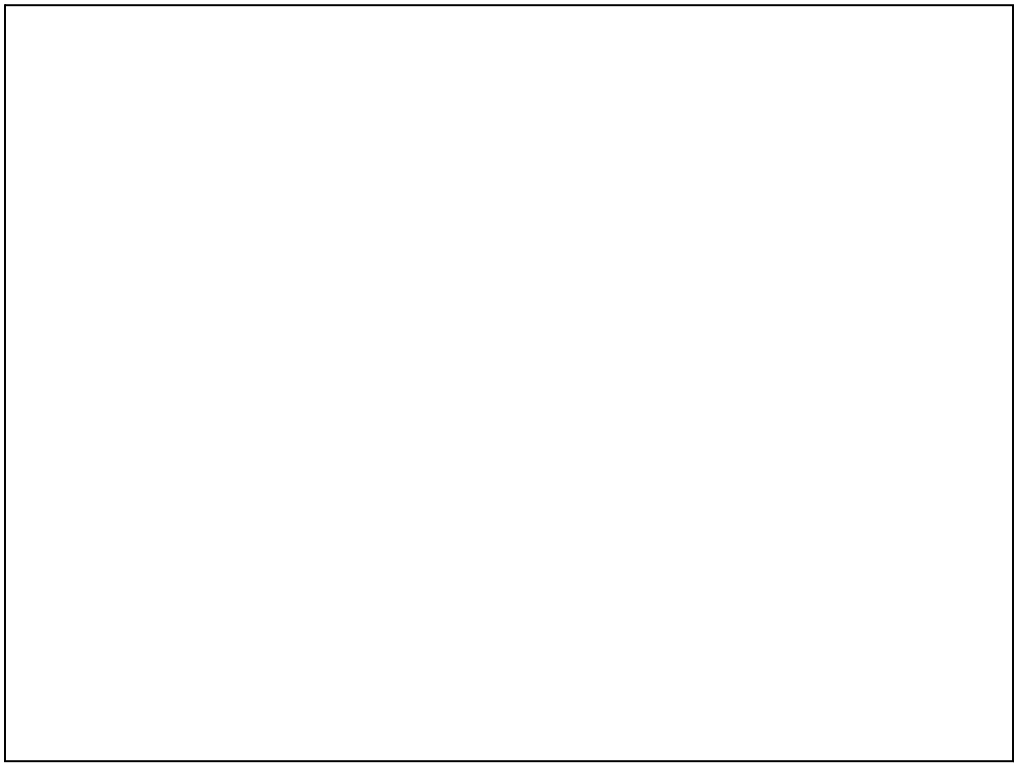

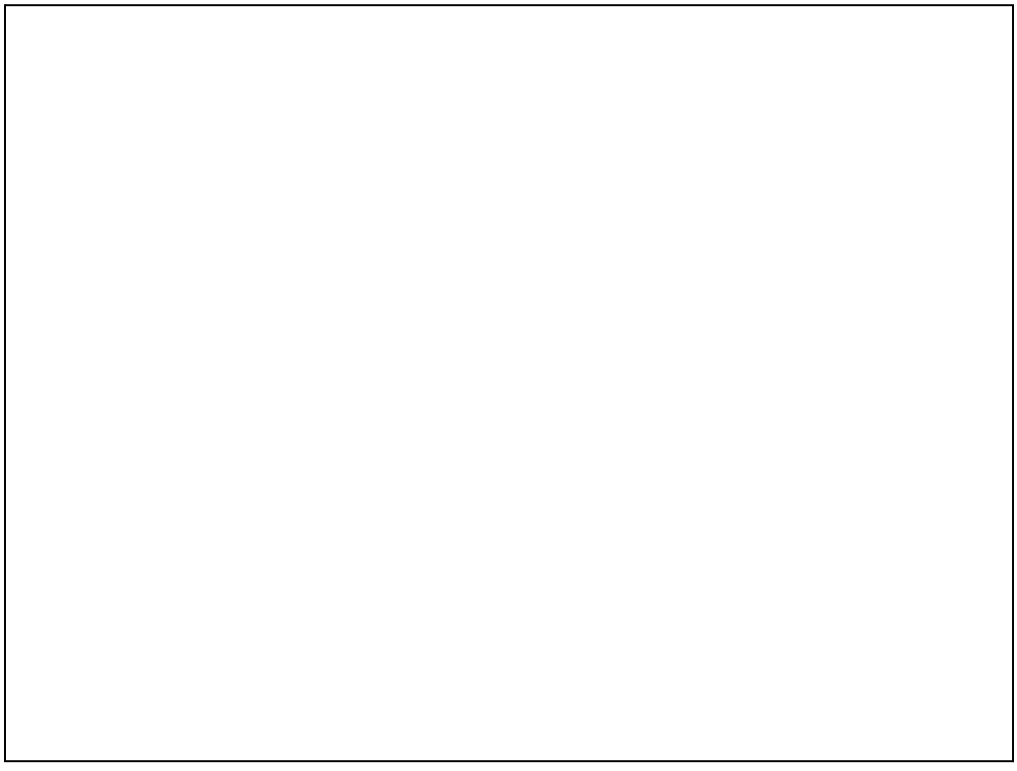

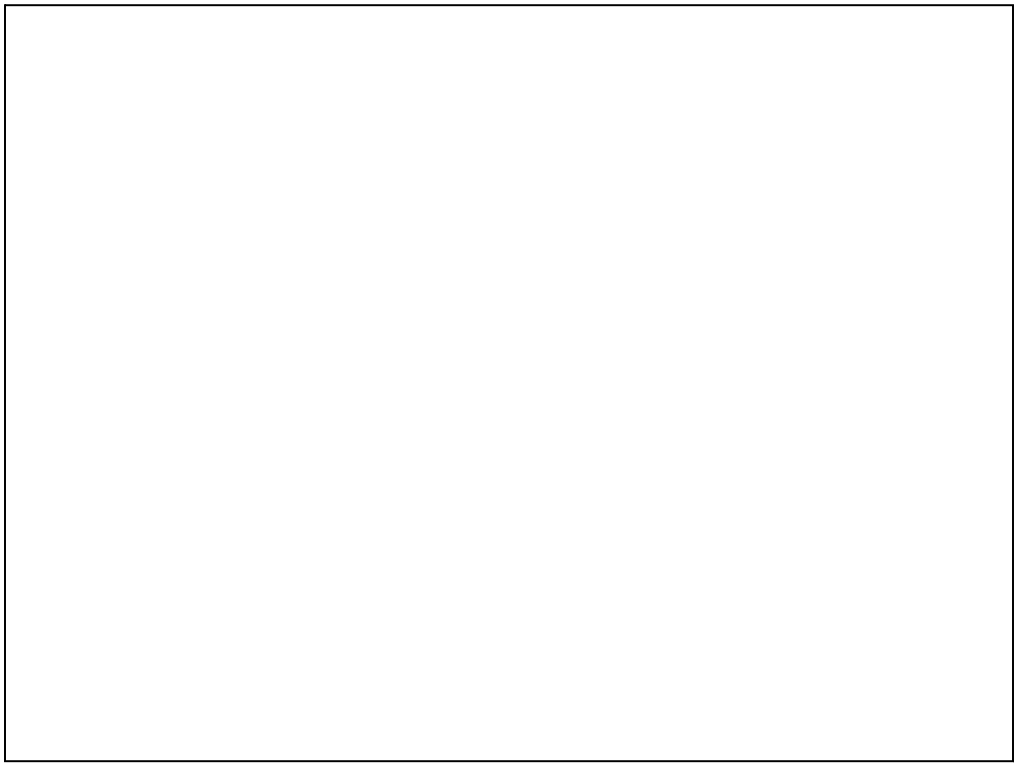

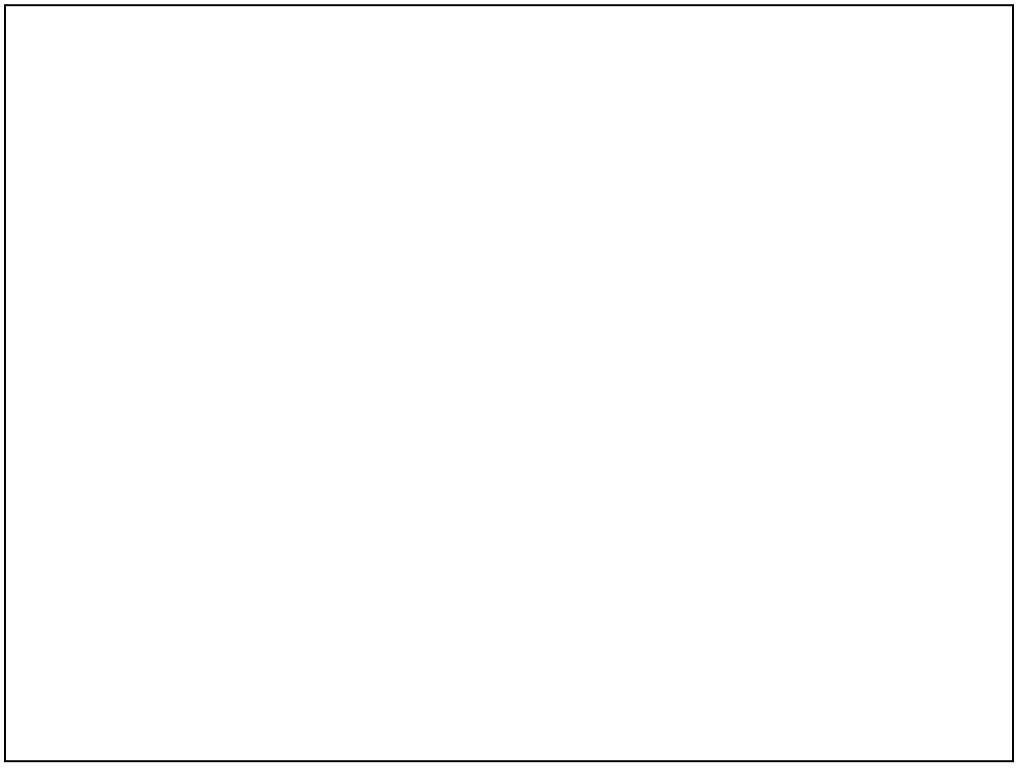

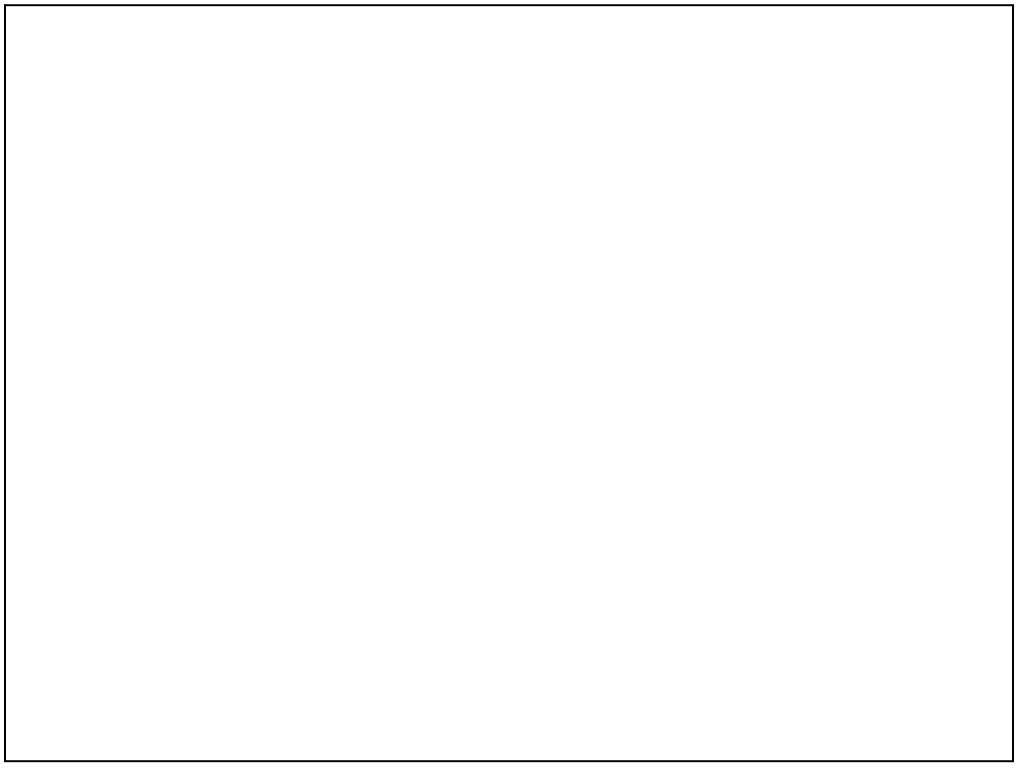

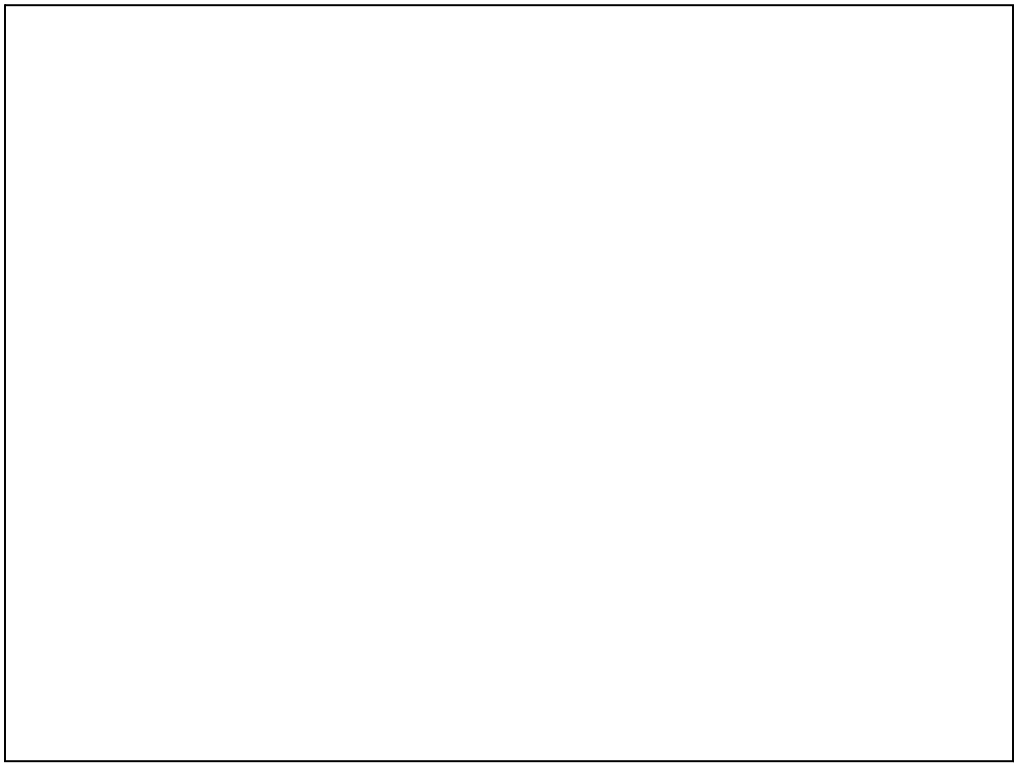

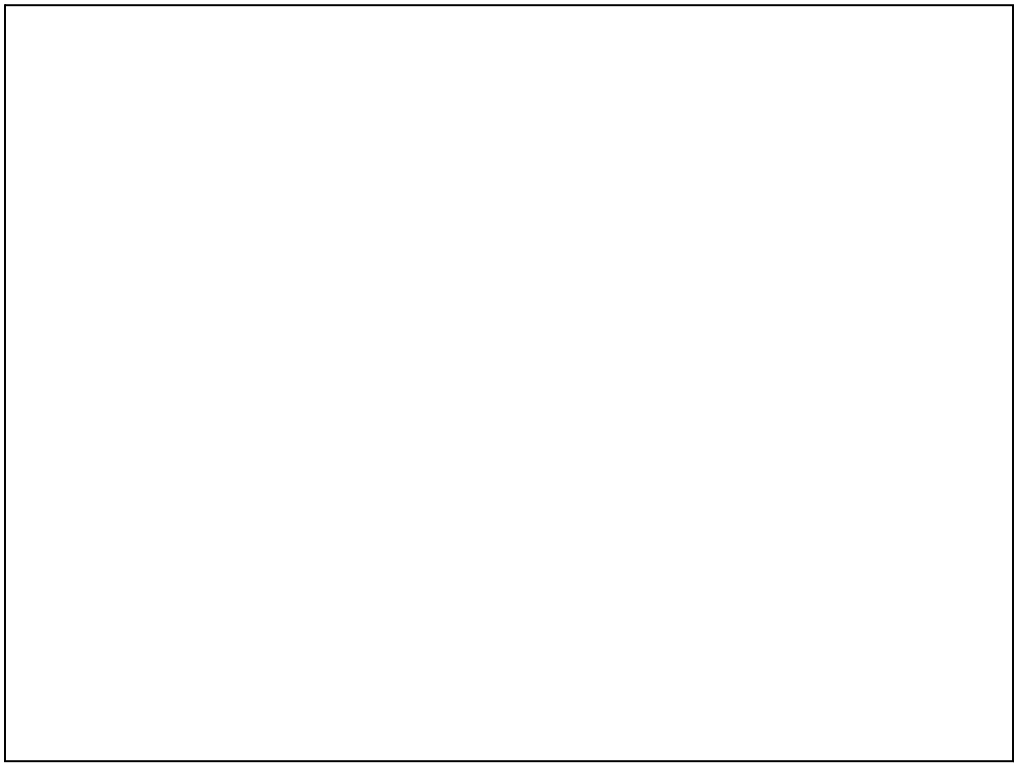

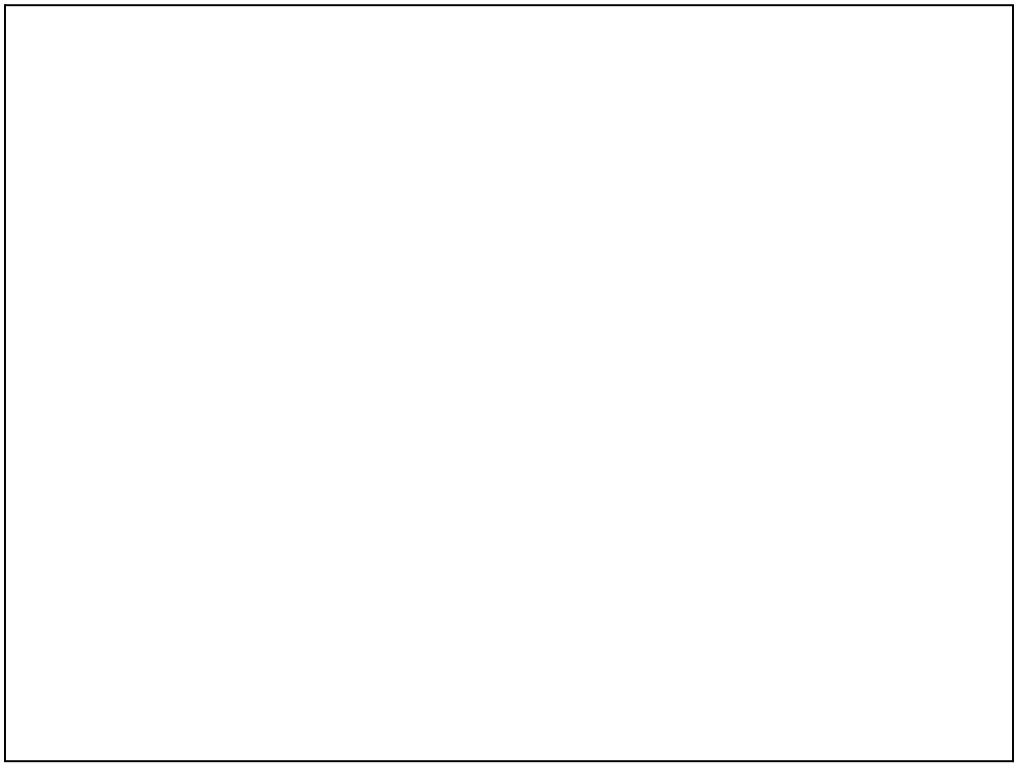

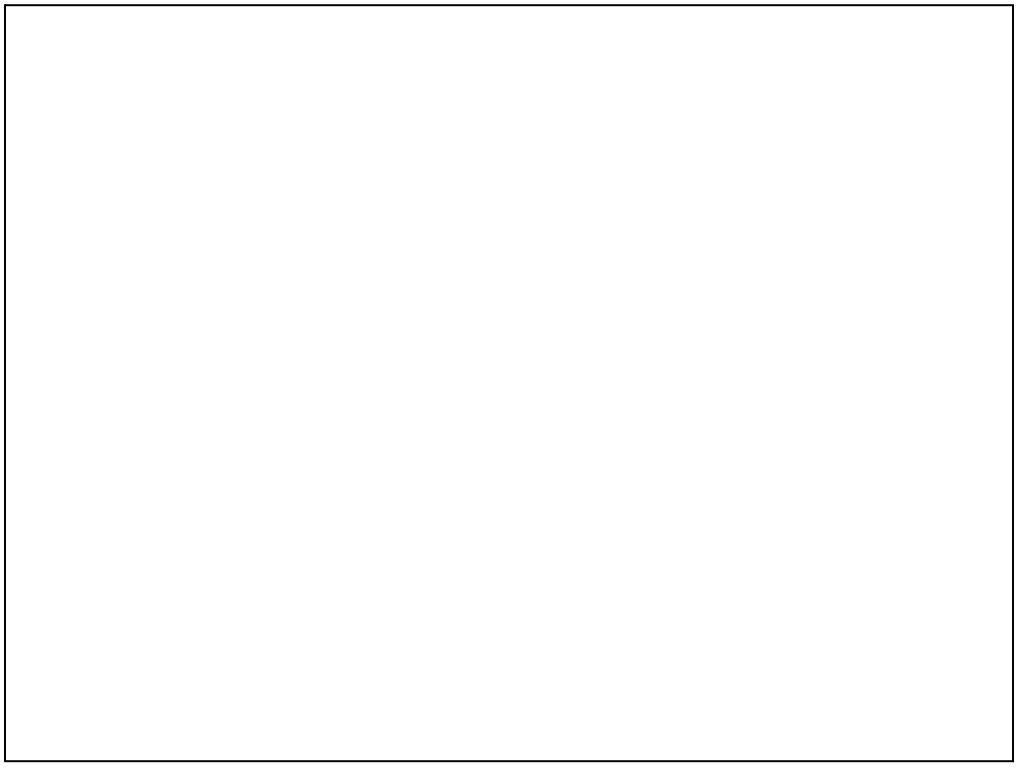

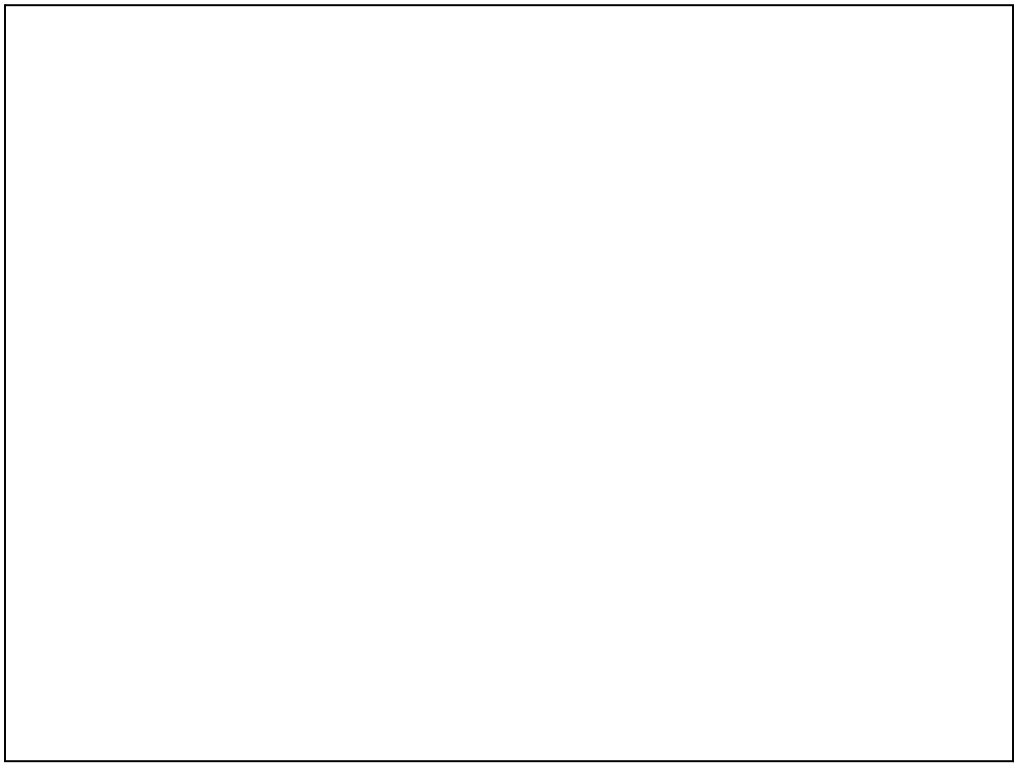

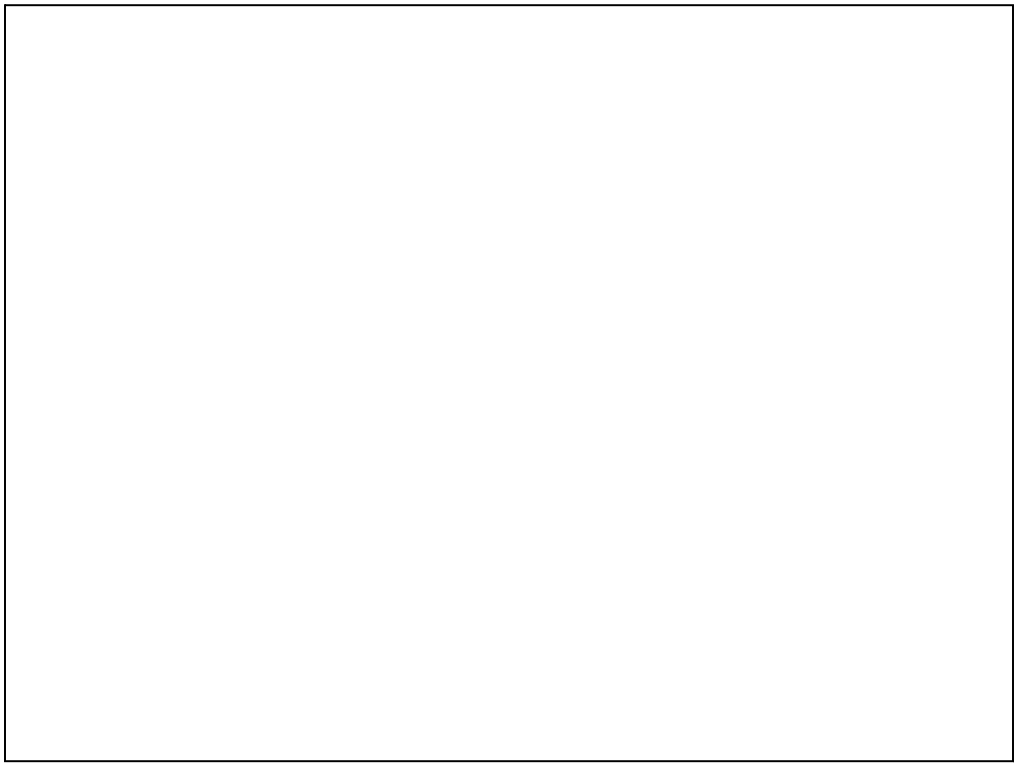

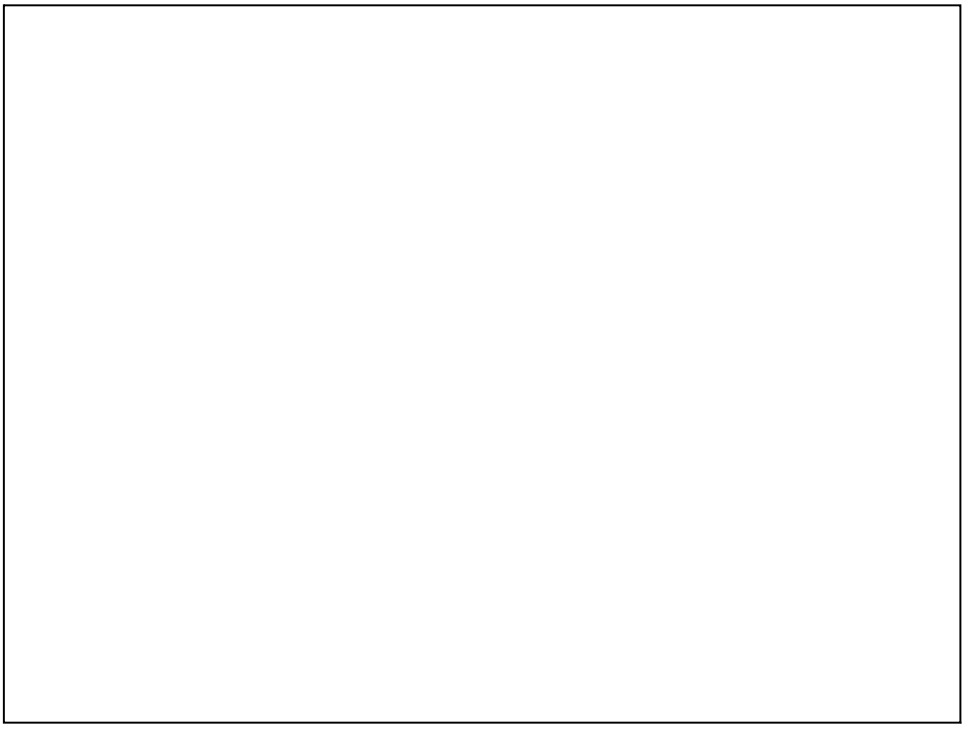

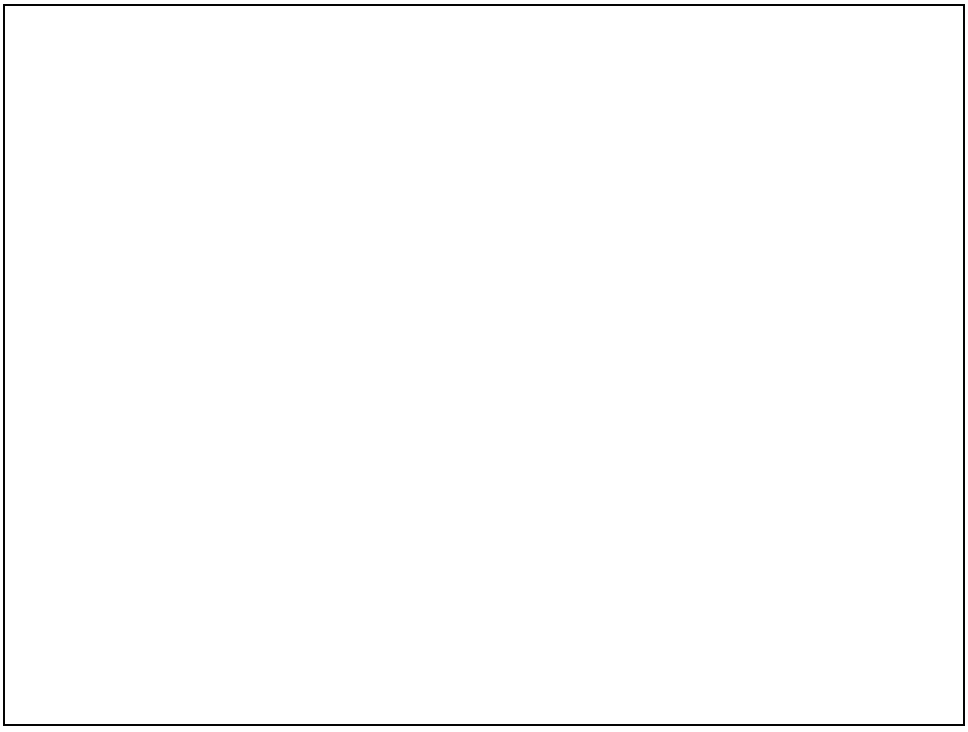

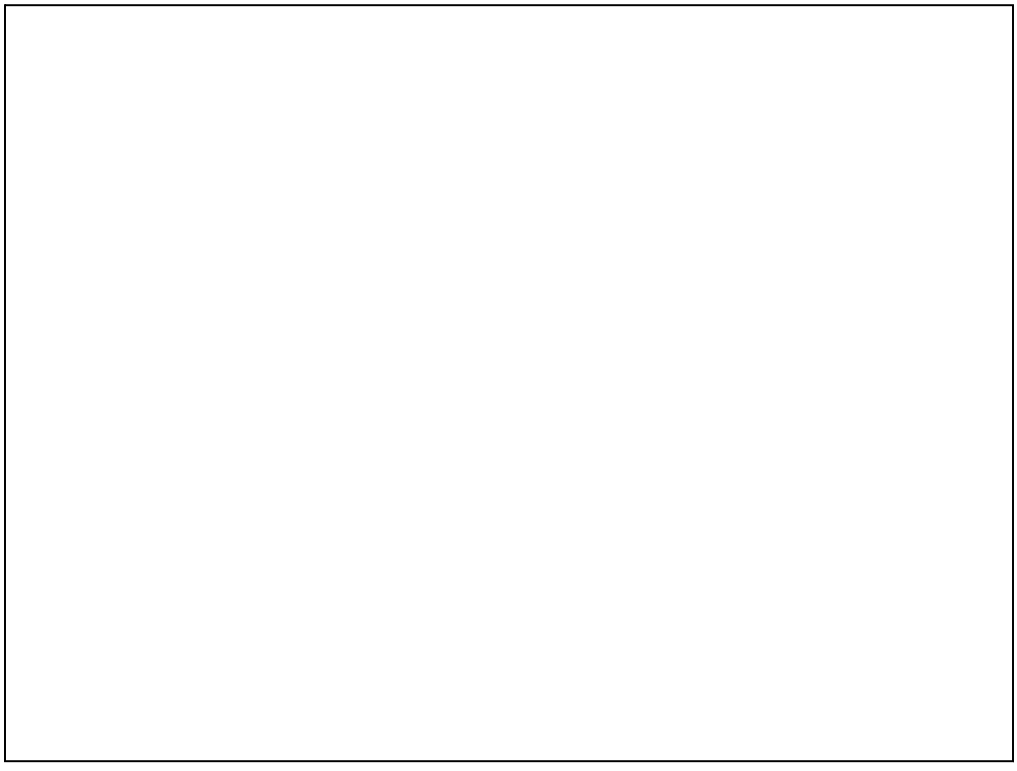

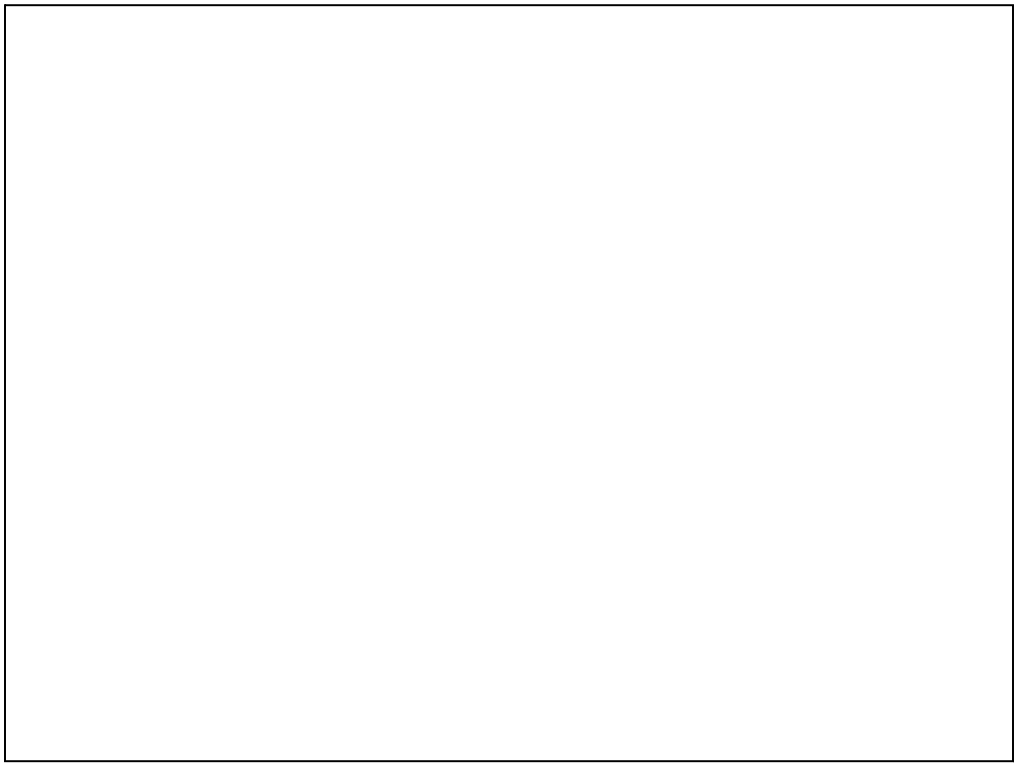

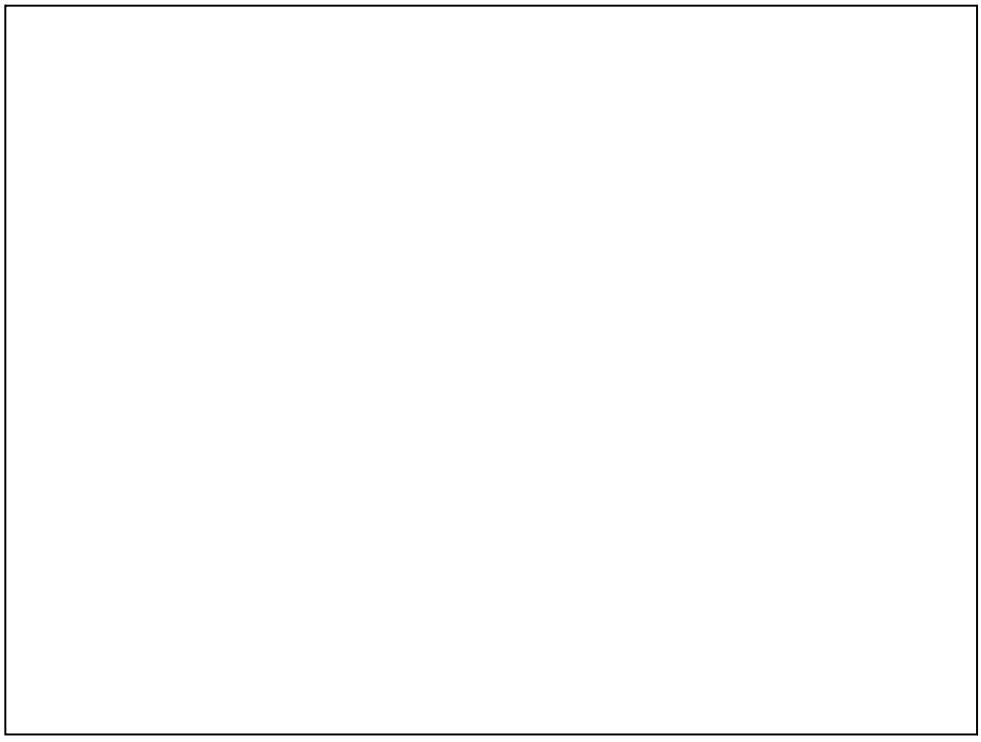

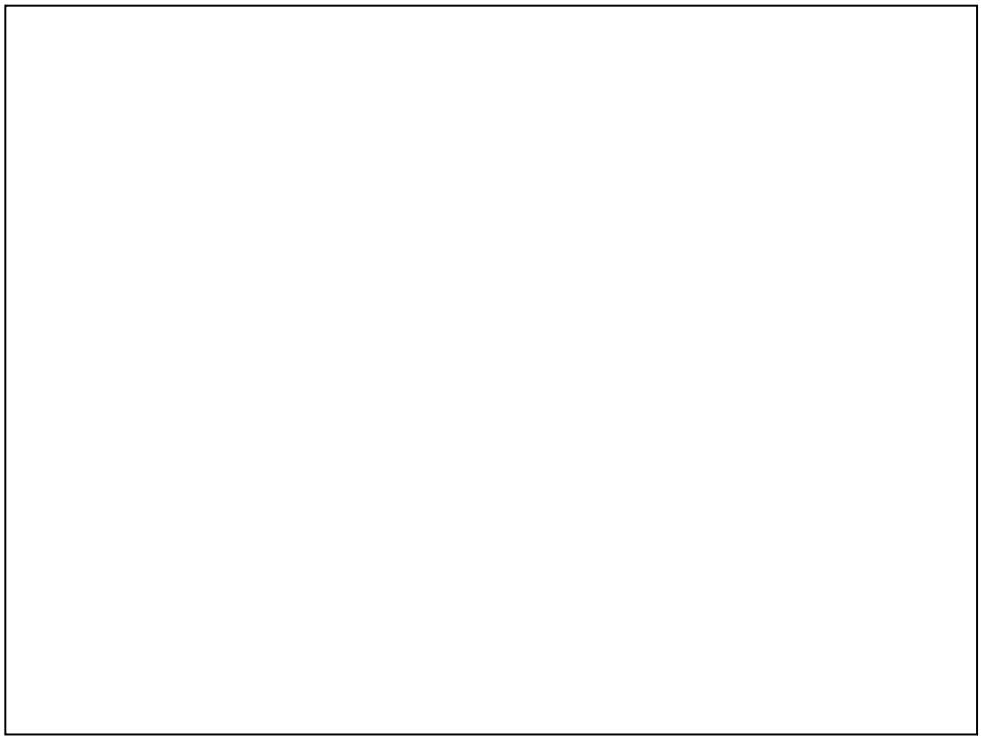

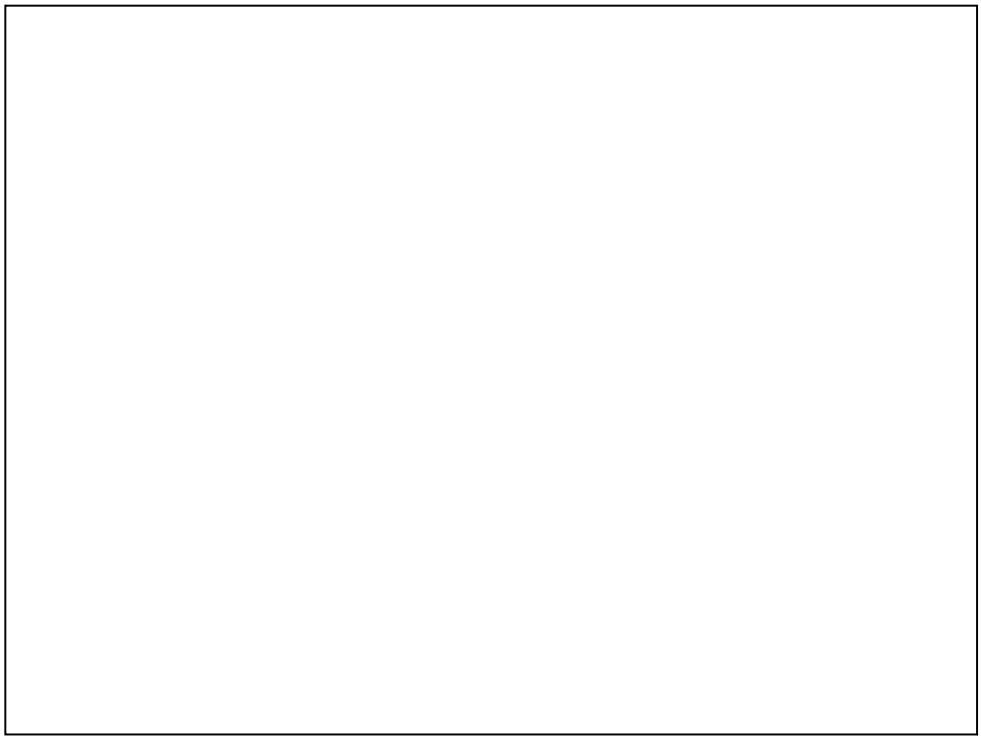

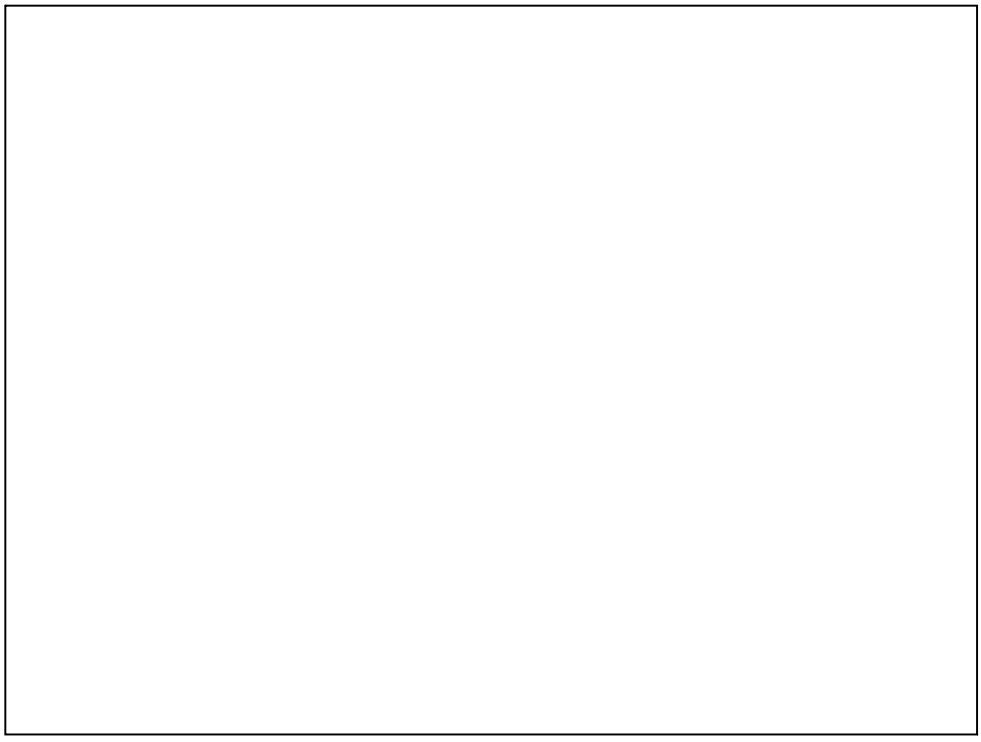

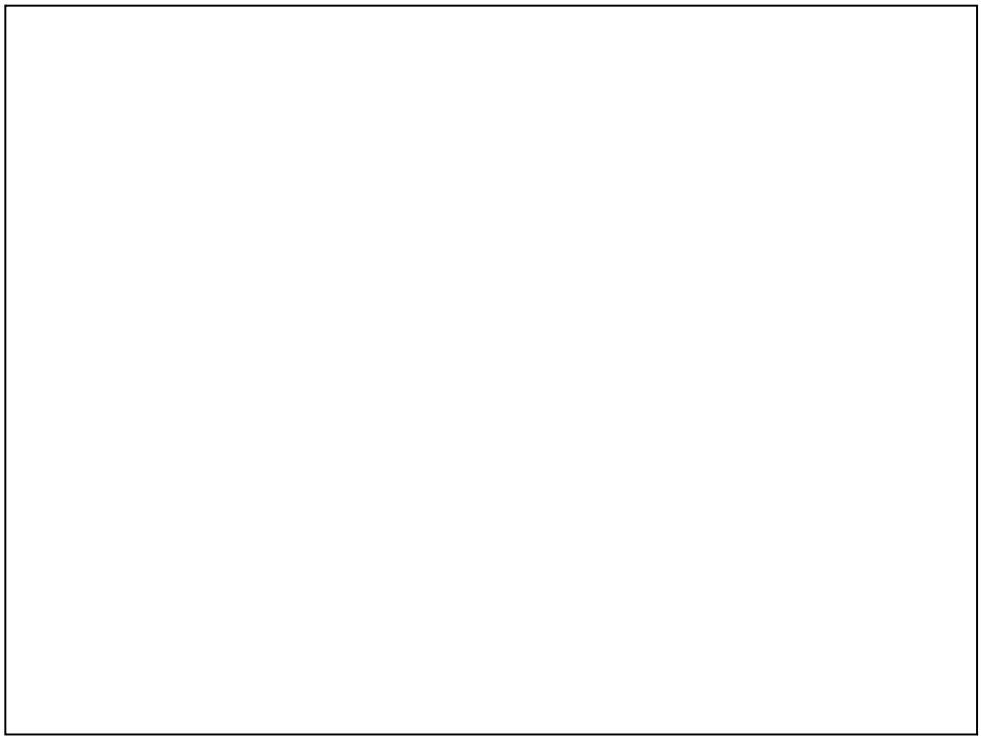

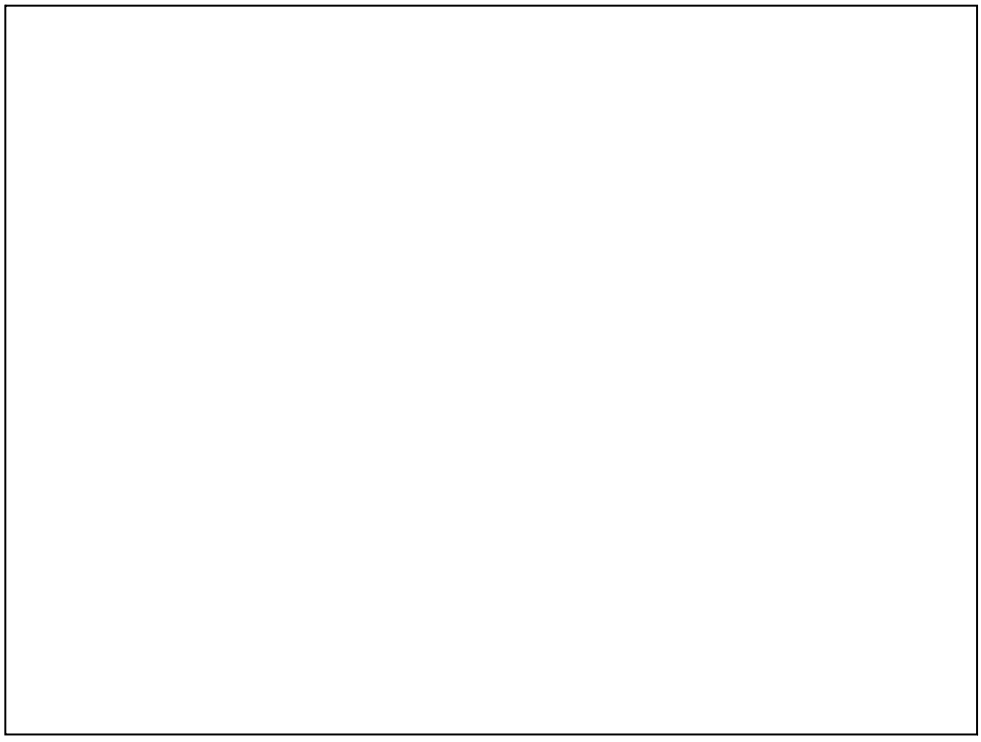

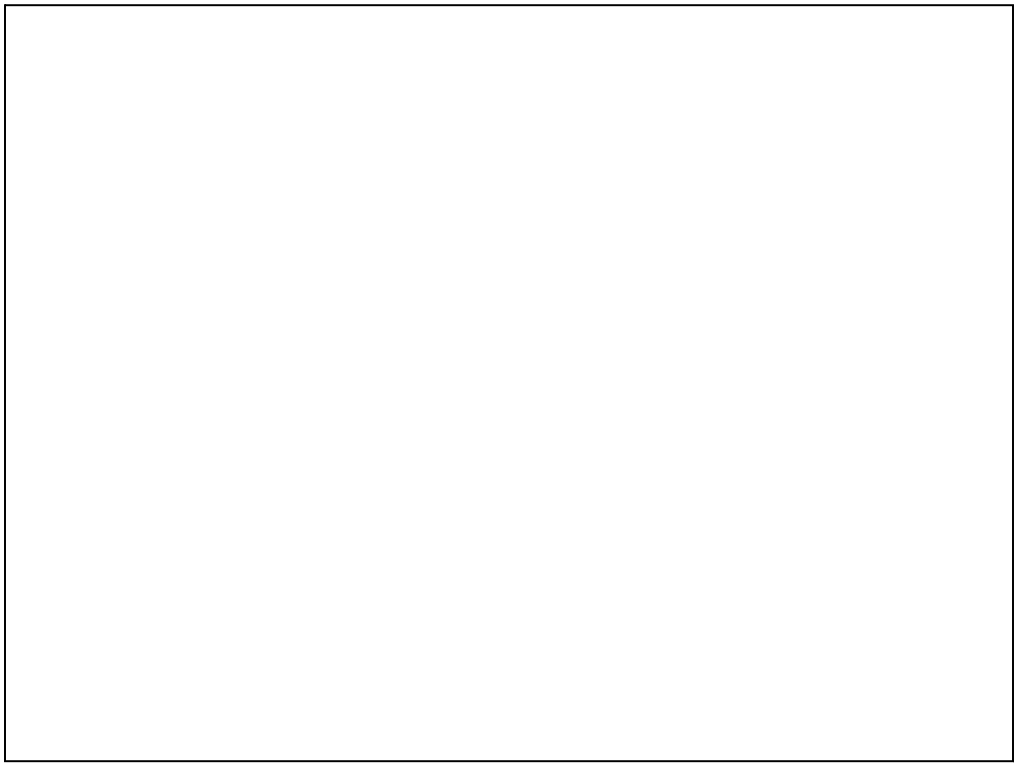

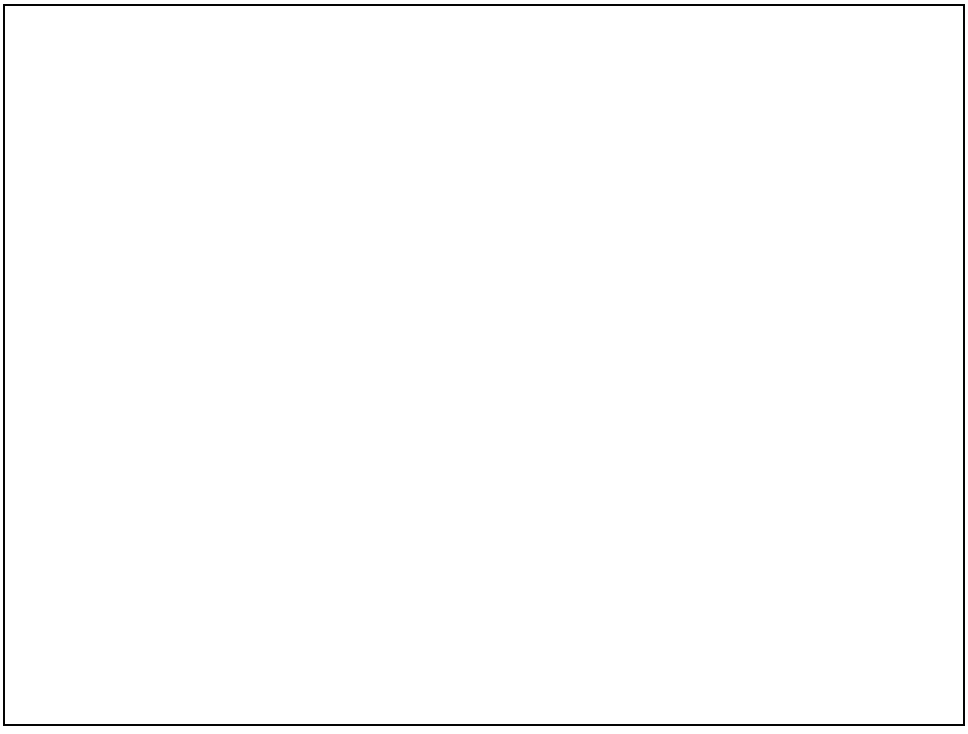

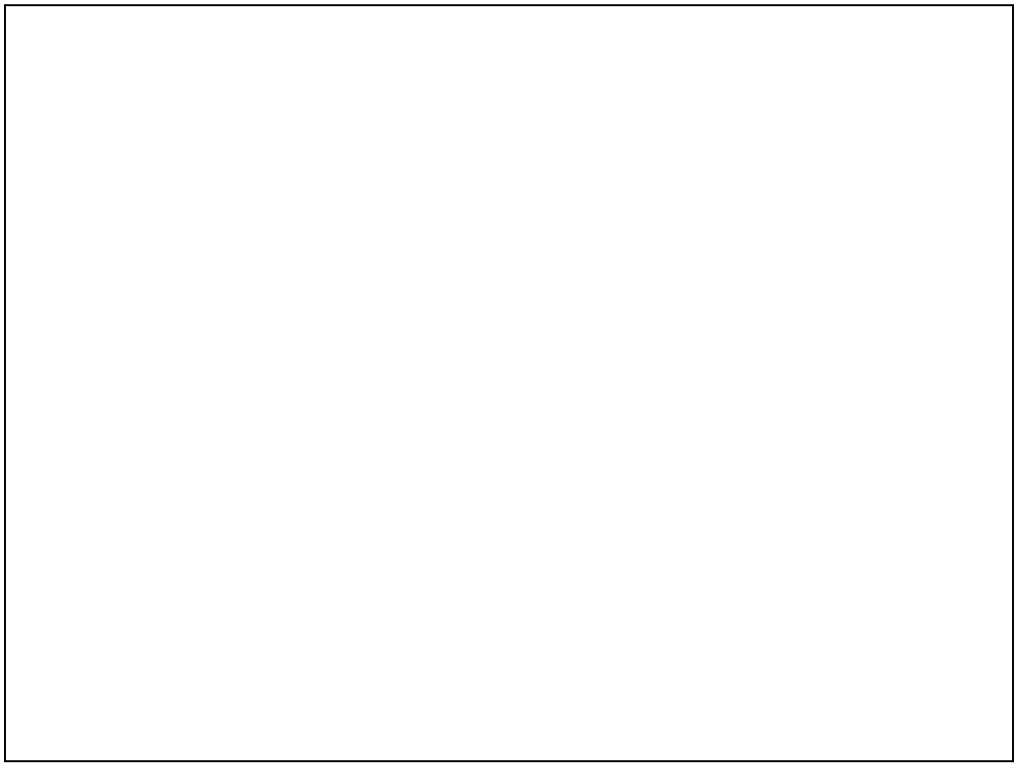

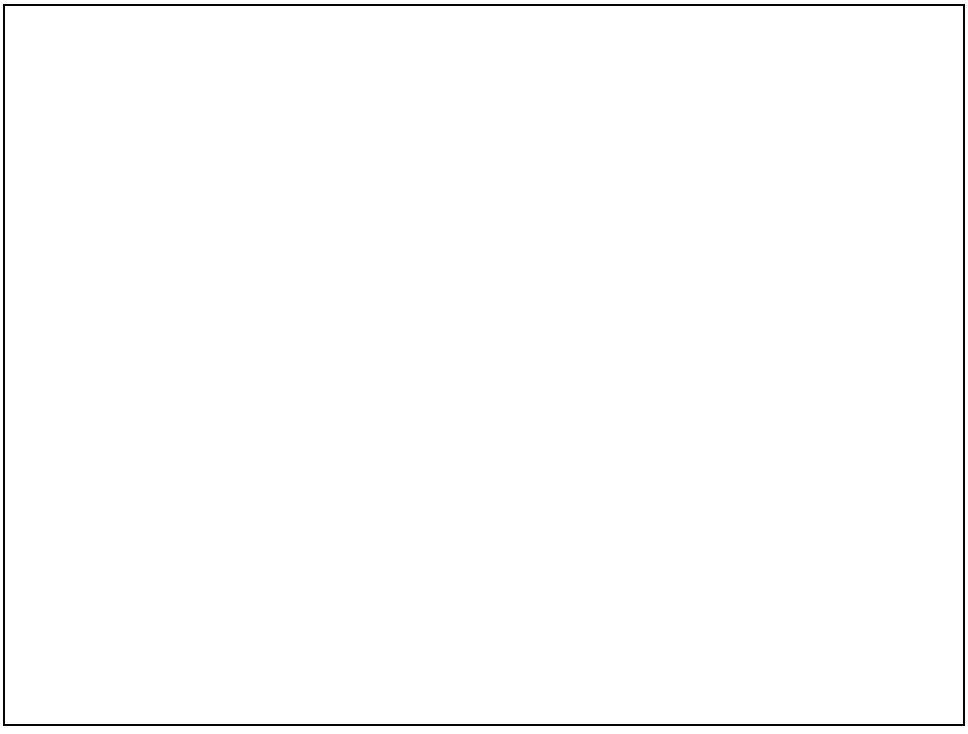

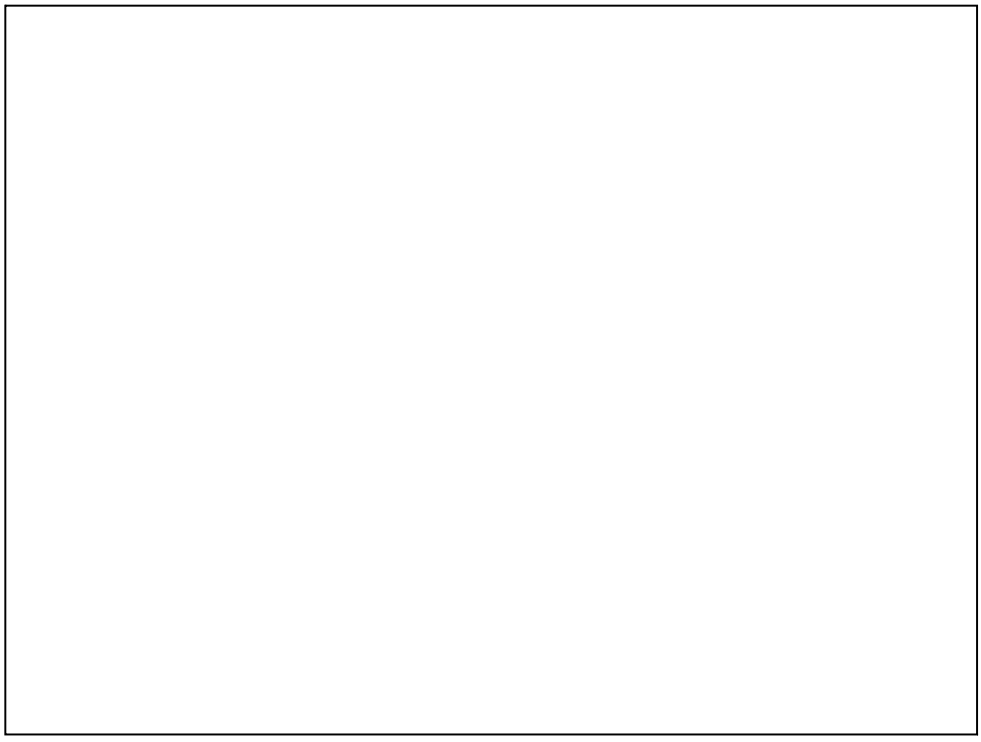

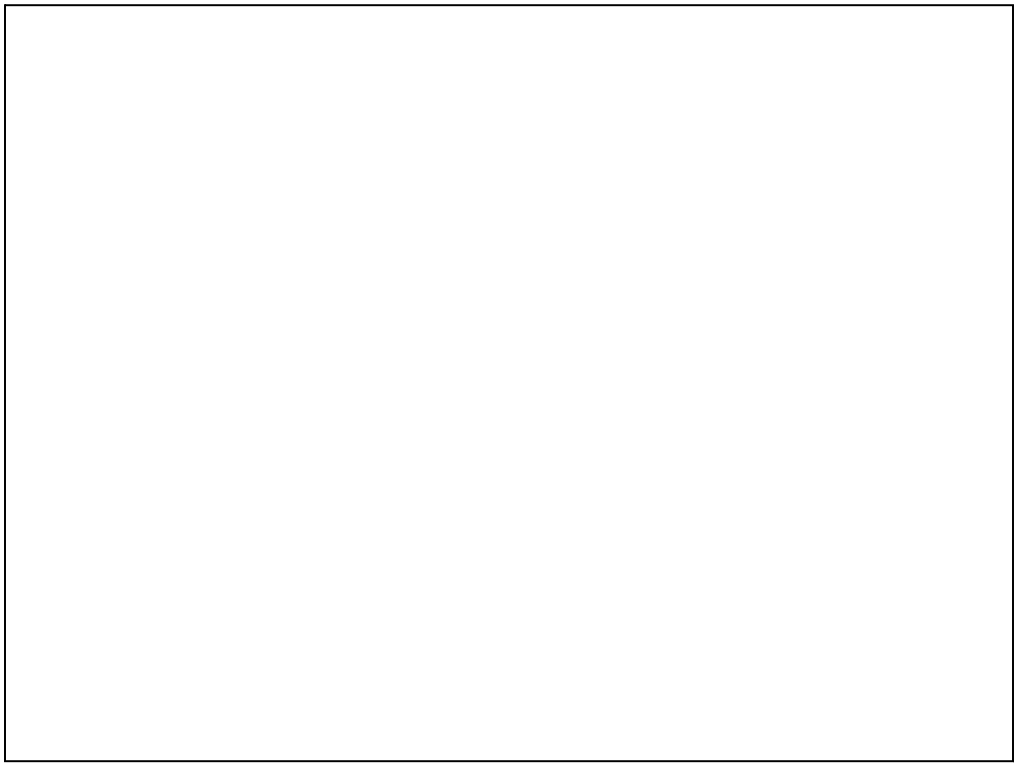

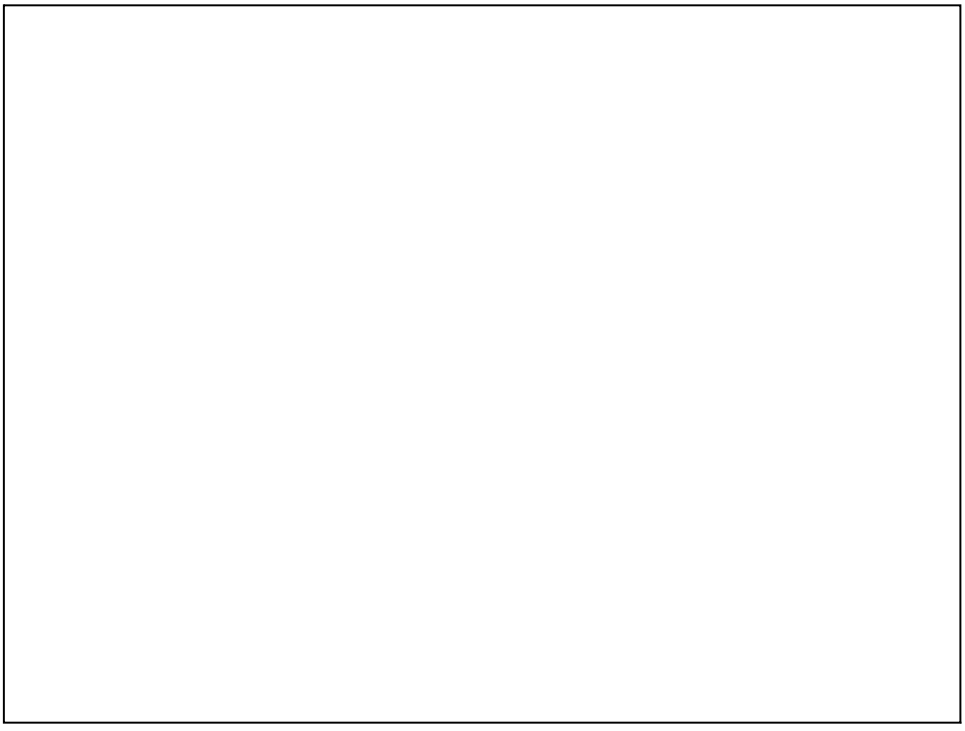

Which path will RIPE pick?

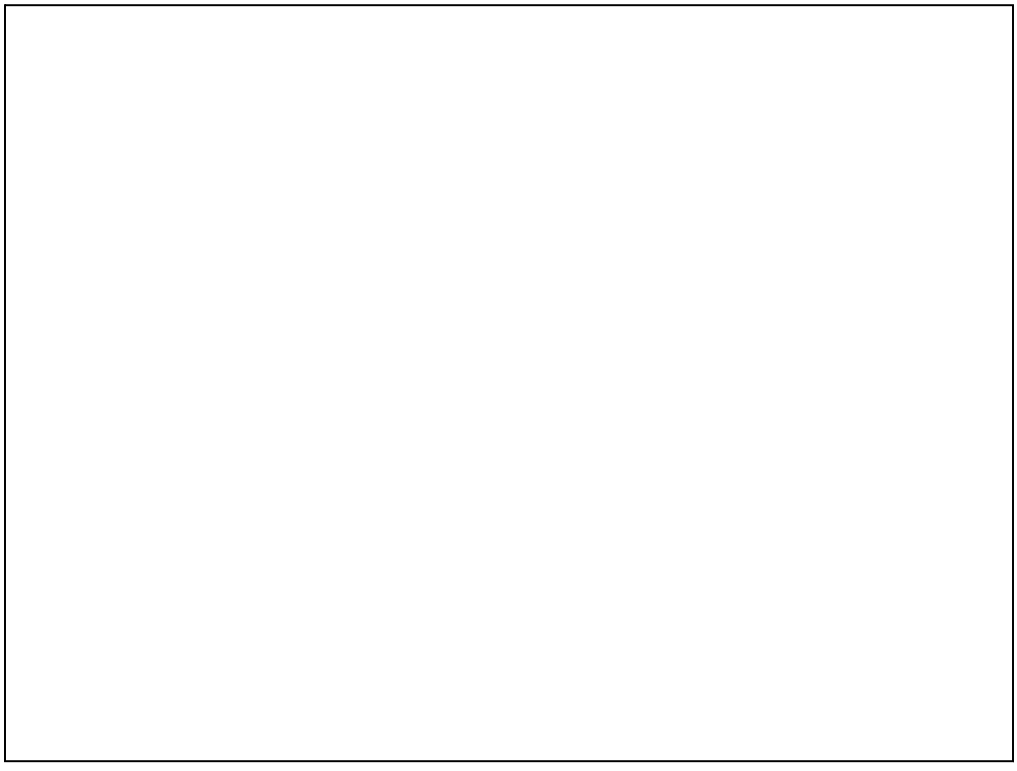

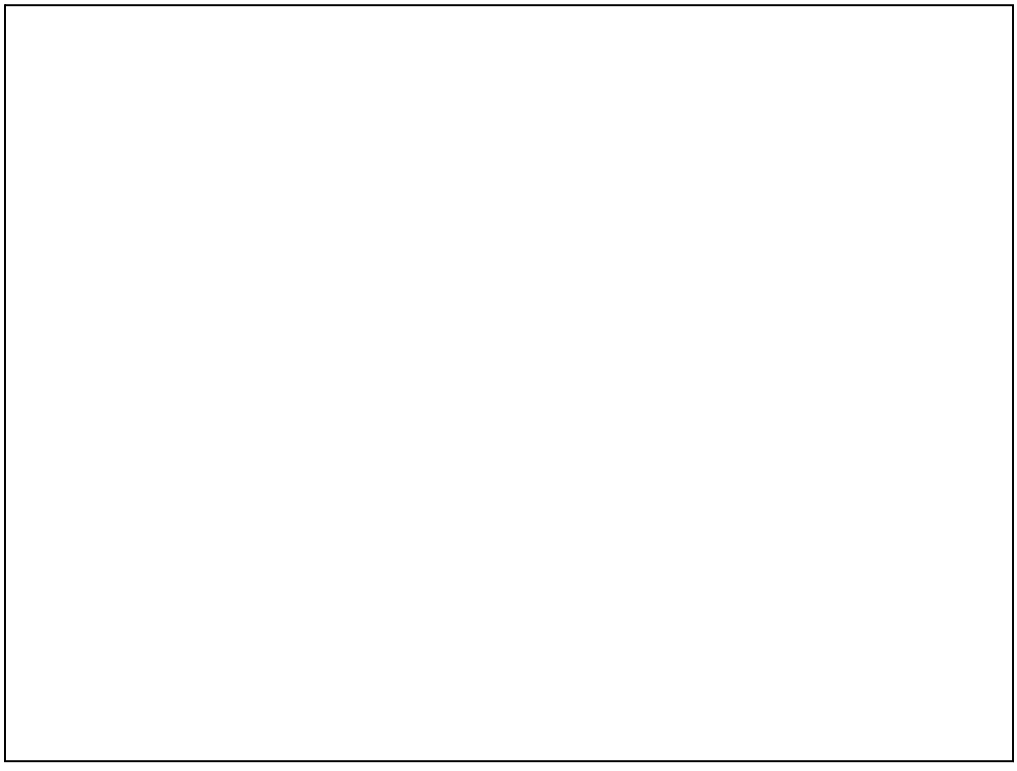

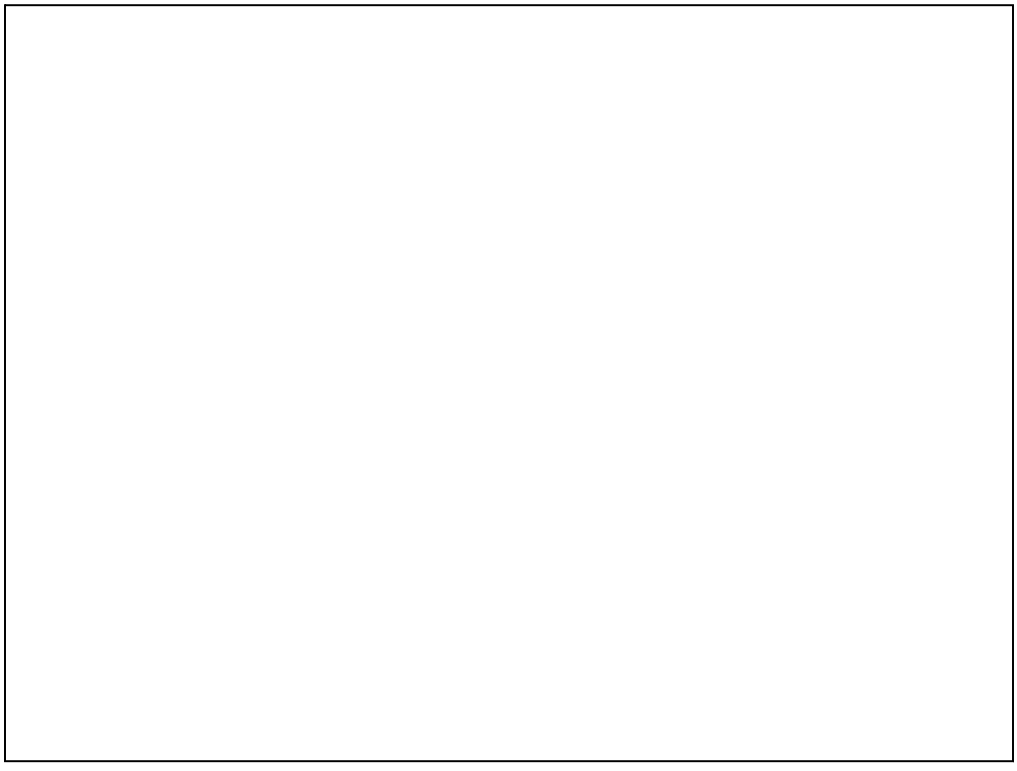

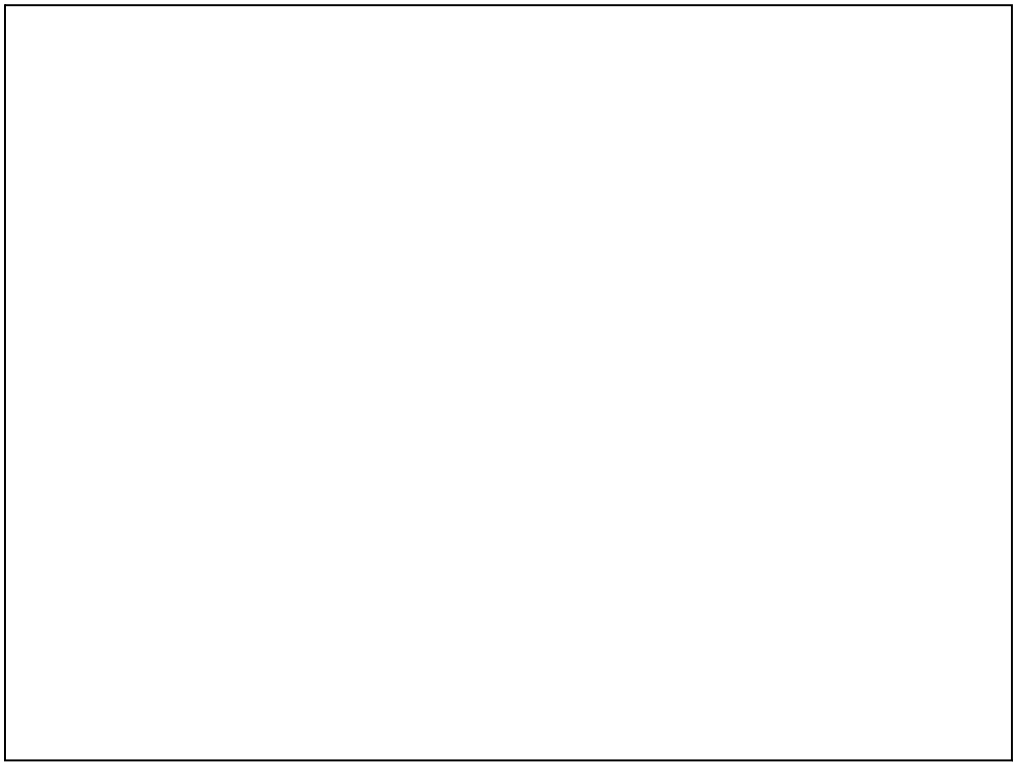

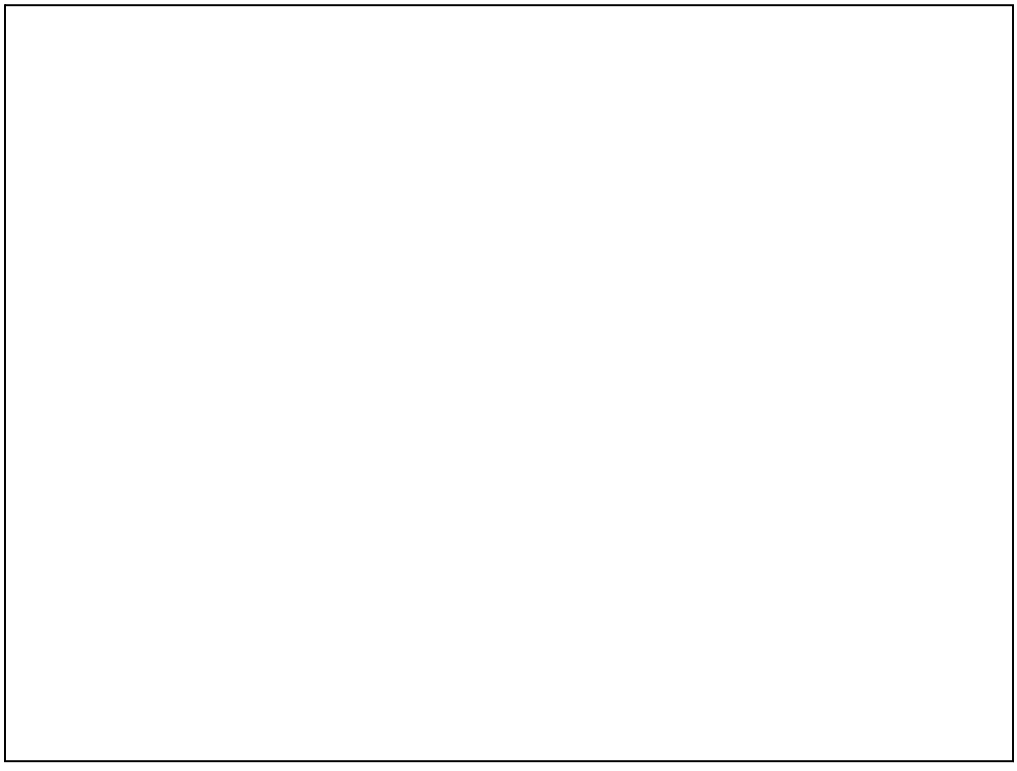

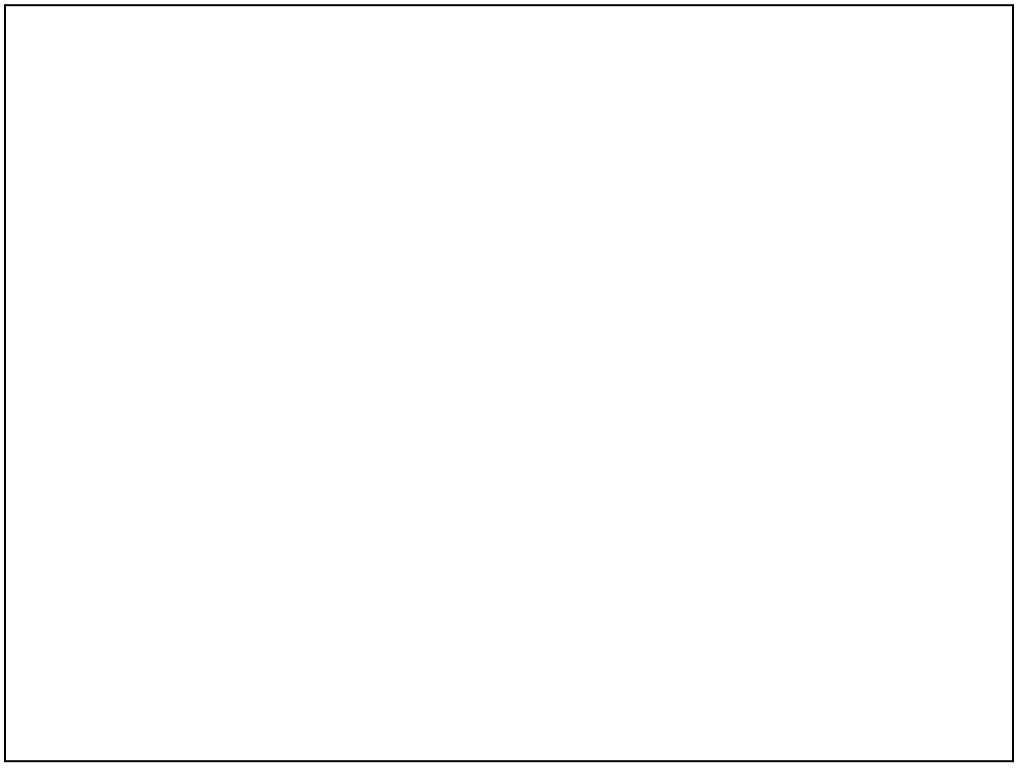

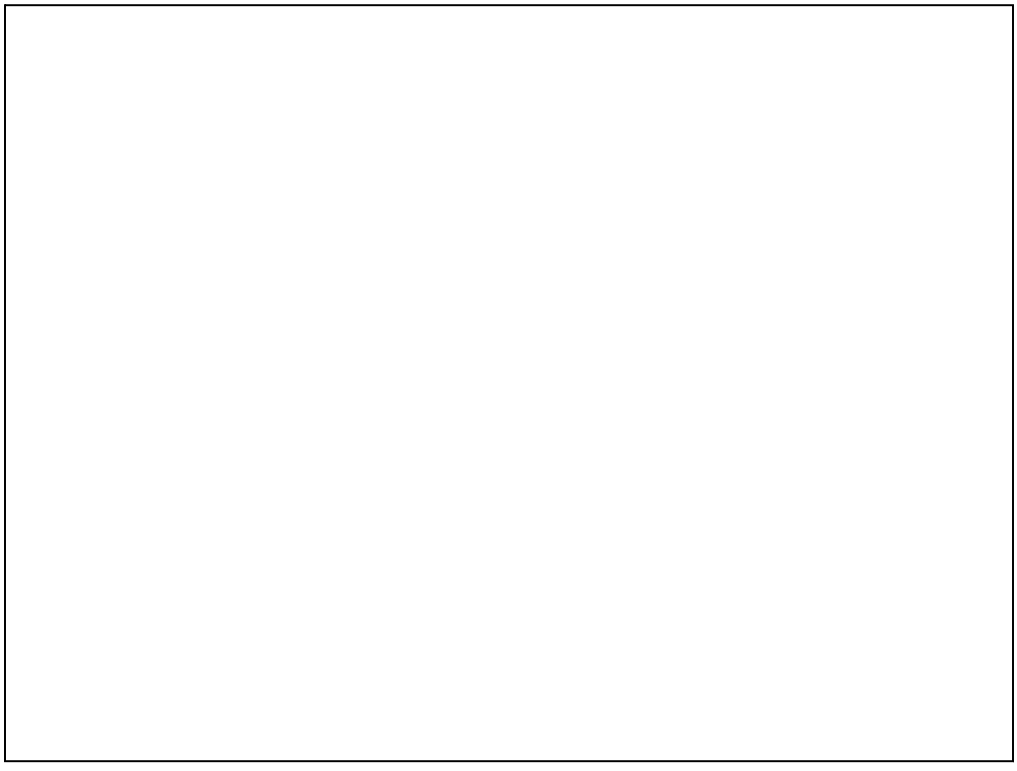

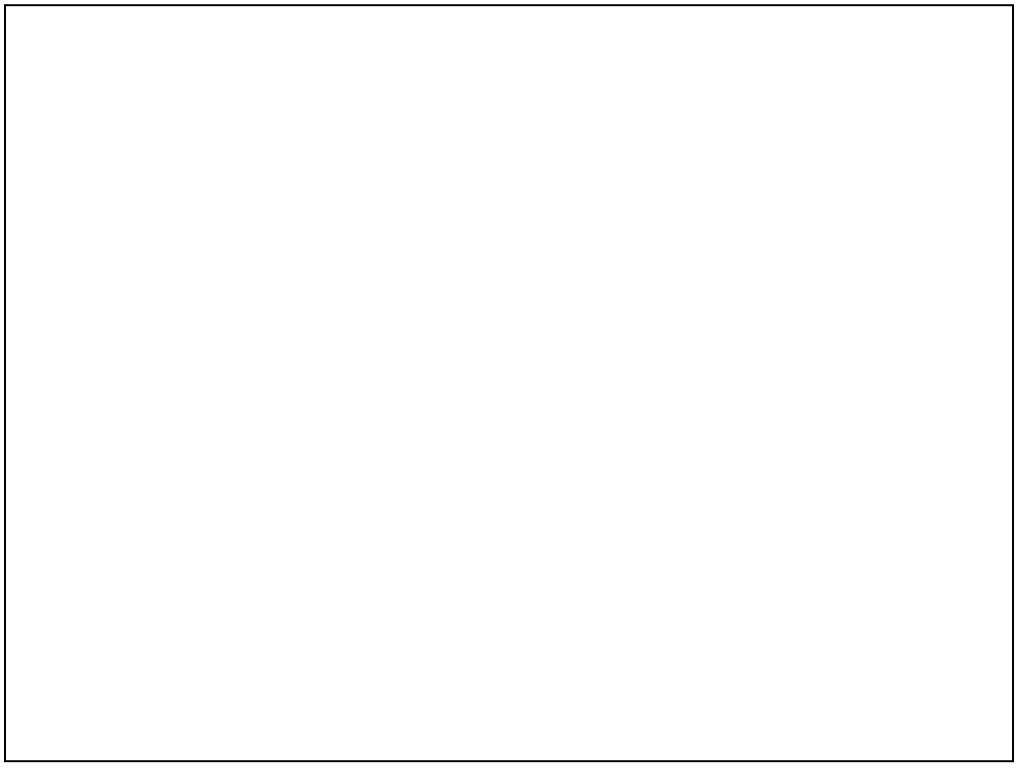

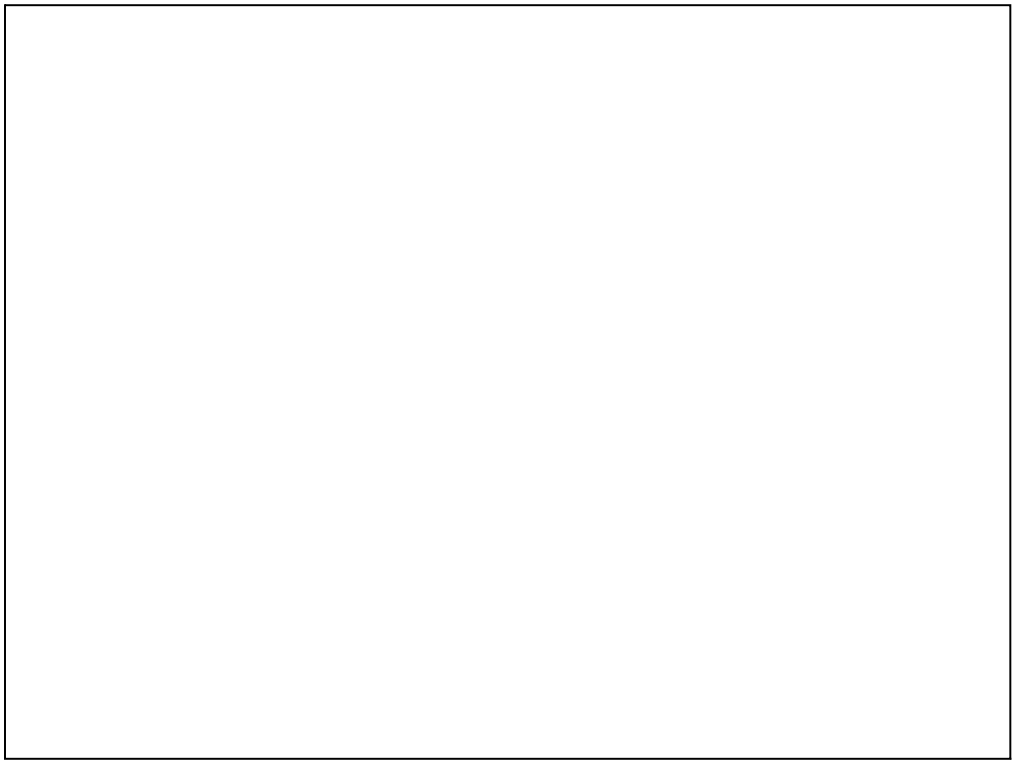

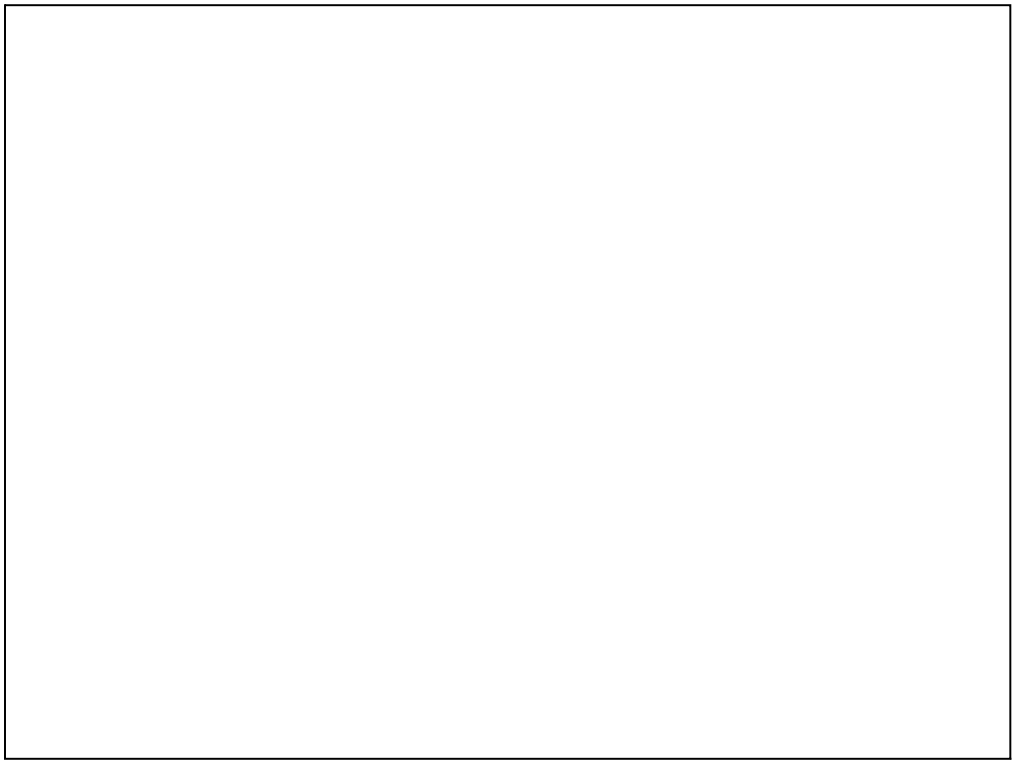

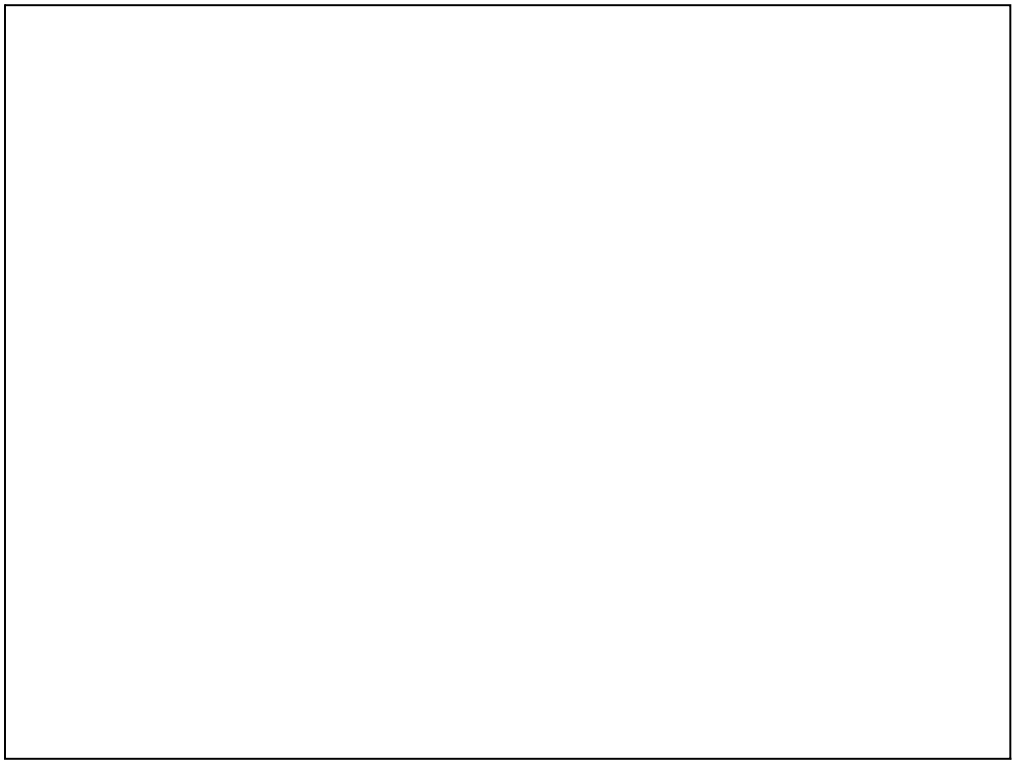

BGP Security Problems

BGP Security Problems

- Infrastructure built on top of BGP is highly vulnerable to malicious attacks
- Fundamentally, no way to guarantee BGP router uses allocated AS number or that it holds address space it advertises
- Most of BGP's security problems are result of lack of verification that IP address space belongs to a given AS

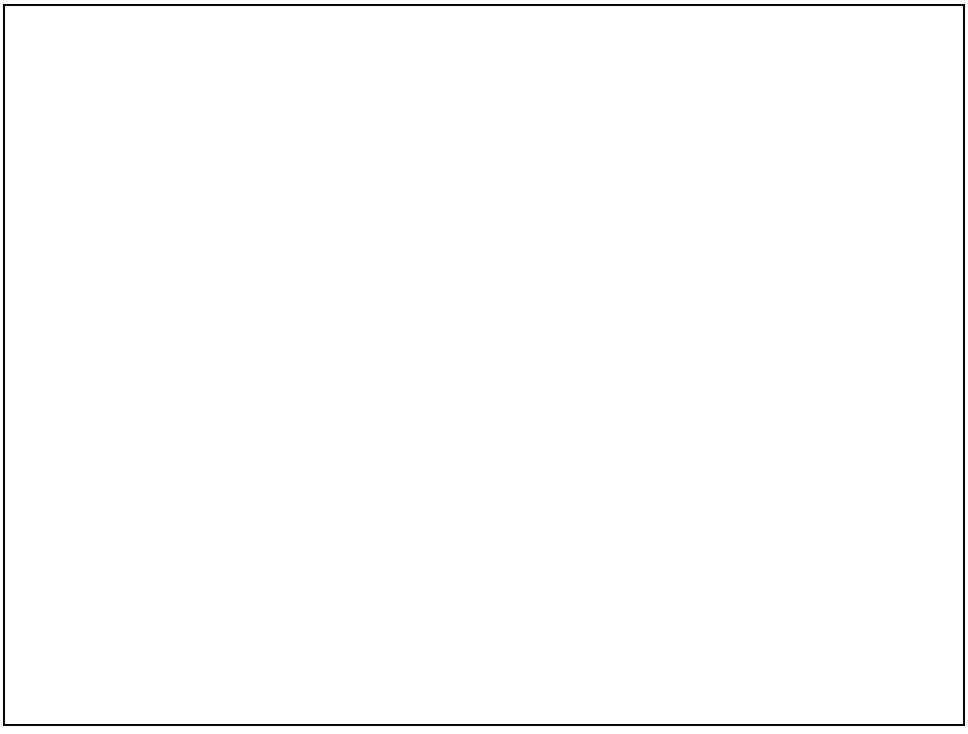

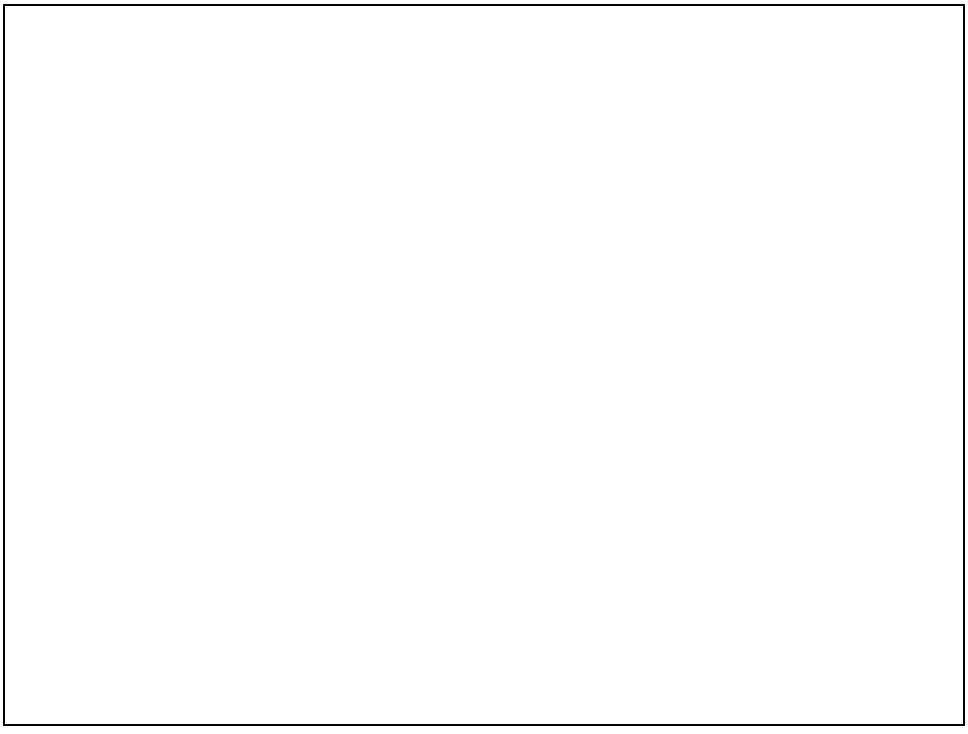

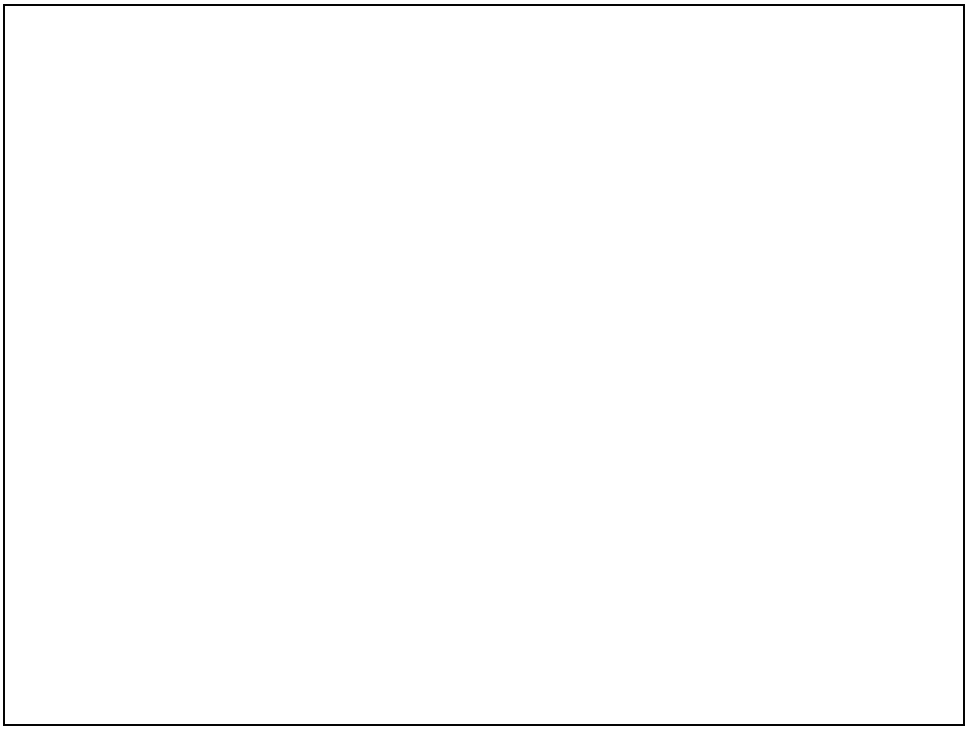

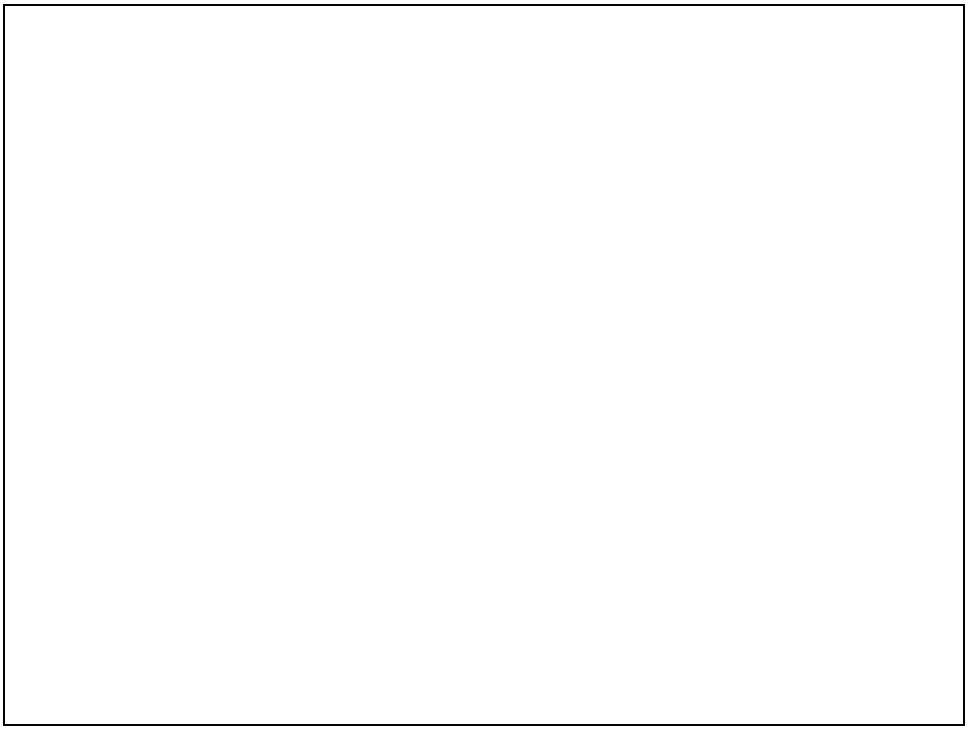

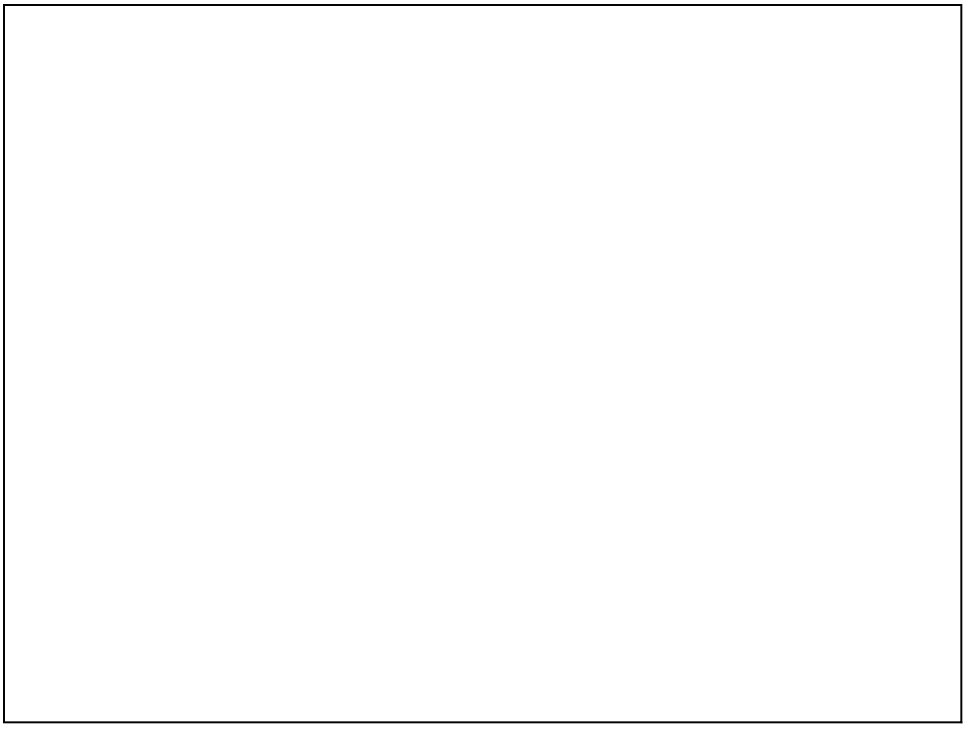

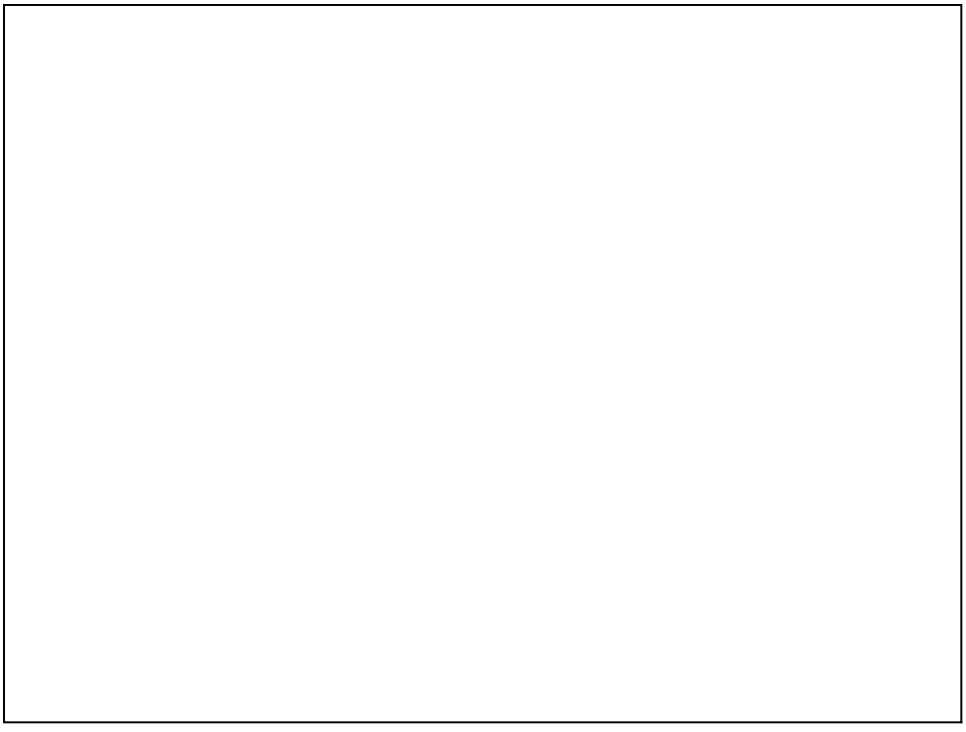

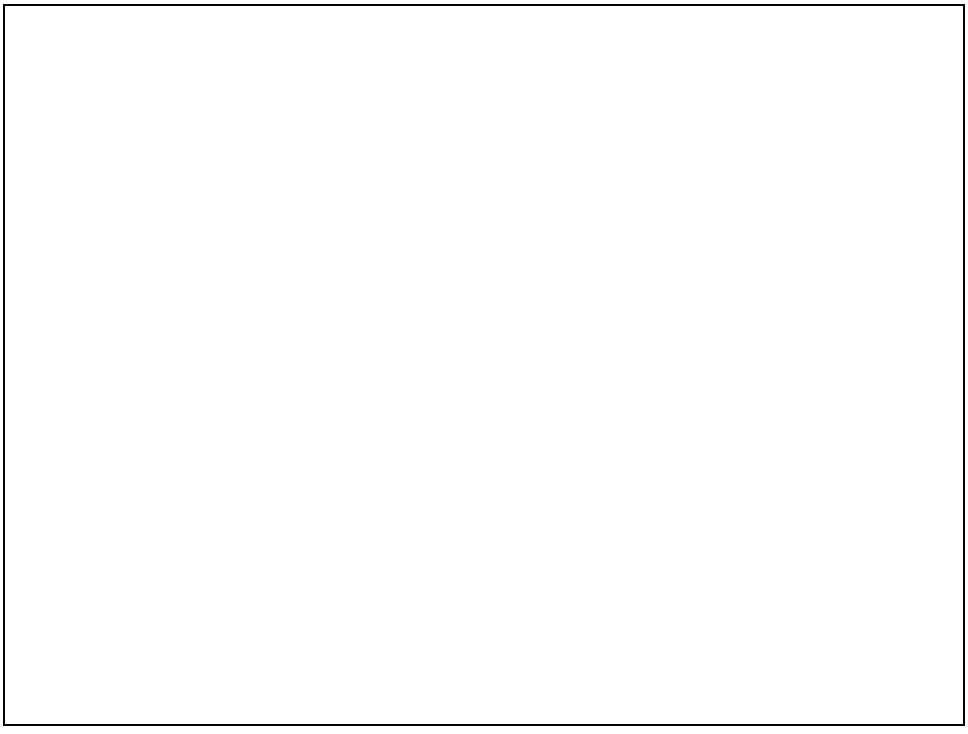

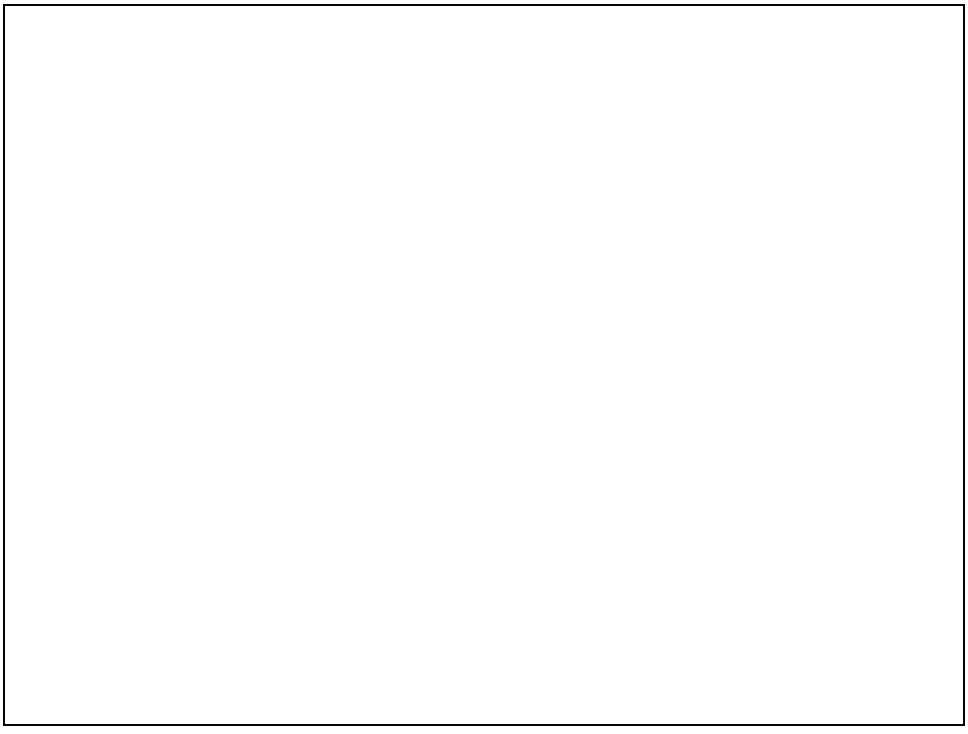

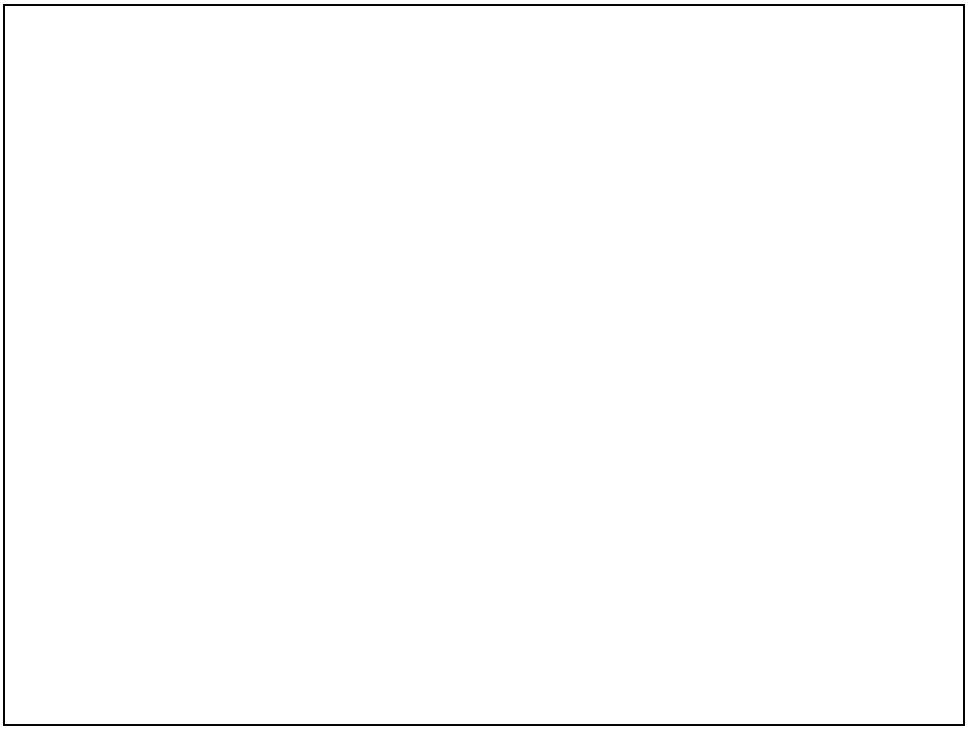

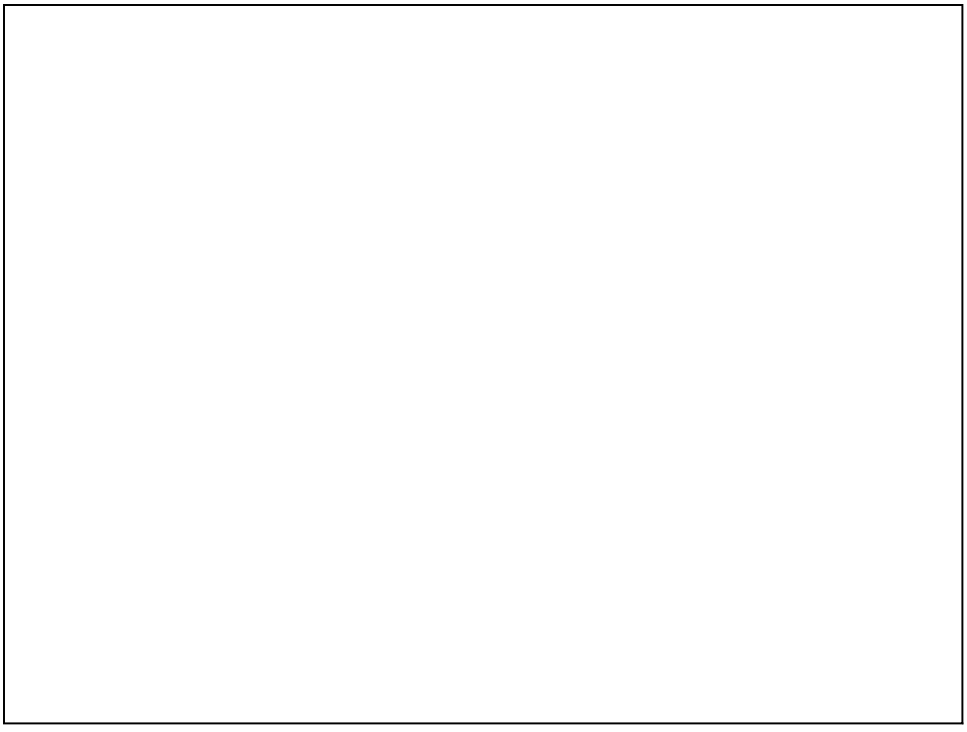

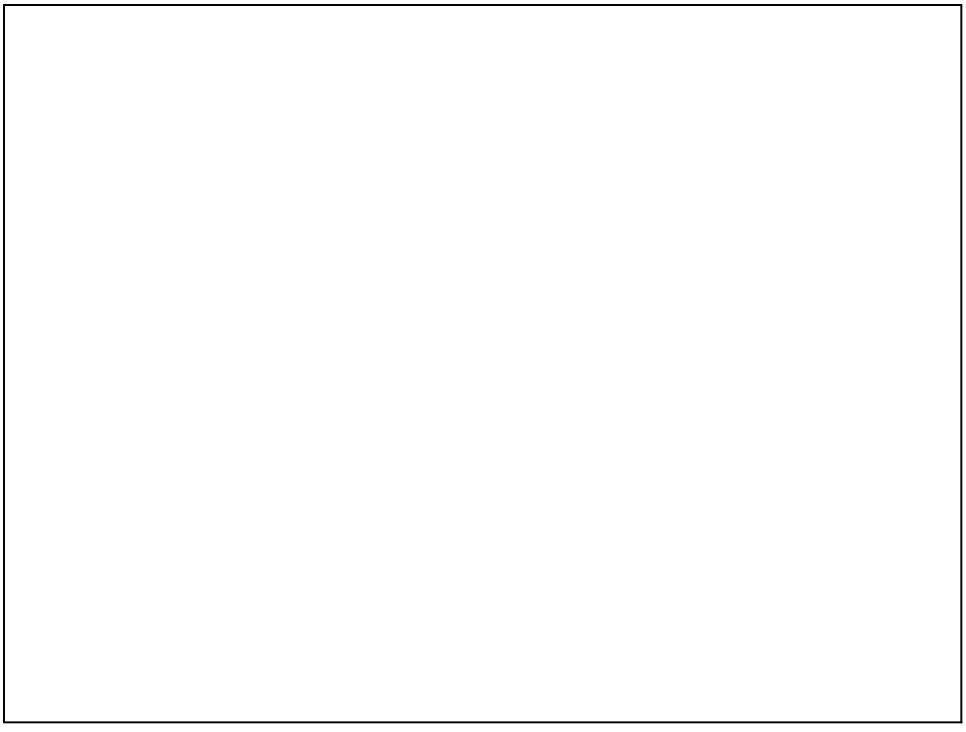

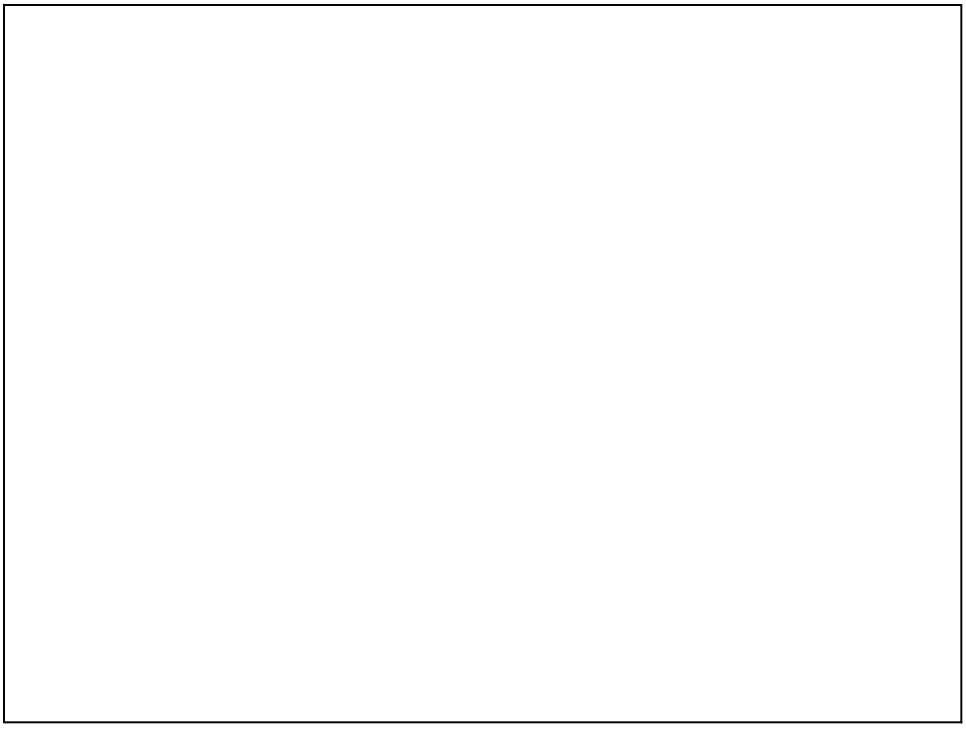

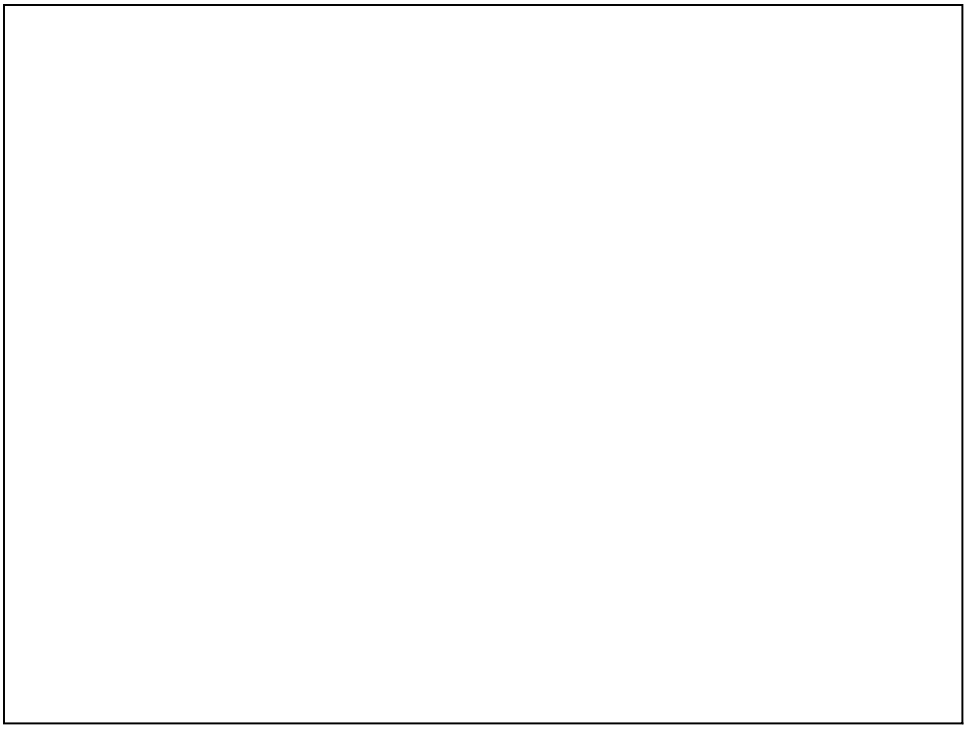

Current BGP Attacks

- **2013 Renesys, Provides Internet Intelligence**
- February 2013, observed sequence of events, lasting from few minutes to several hours in duration, in which global traffic was redirected to Belarusian ISP GlobalOneBel
    - These redirections took place on an almost daily basis throughout February
- Affected countries included the US, South Korea, Germany, the Czech Republic, Lithuania, Libya, and Iran
- Another set of traffic hijack incidents took traffic to Iceland

http://www.renesys.com/2013/11/mitm-internet-hijacking/