

BGP/OSPF Implementation

Link Technologies, Inc.

Dennis Burgess

- **Mikrotik Certified Trainer / Engineer**
- **MikroTik Certified Dude Consultant**
- **Consulting Since 1997**
 - *Enterprise Class Networks*
 - *WAN Connectivity*
- **Certifications**
 - *Cisco, Microsoft, MikroTik*
- **BGP/OSPF Experience**
 - *Deployed many BGP and OSPF networks based on MikroTik, Cisco and Juniper*
- **What I do Currently**
 - *Work with WISPs and CLECs all over the world, designing, and assisting in network configurations including wireless, OSPF, BGP, Traffic Management, Firewalling, and other Network Engineering*

What are we going to Cover?

- **BGP**

- What is it?
- Why do I need it?
- What do I need to get started?
- What does it cost? – with ARIN
- BGP Implementation
 - *Things to know!*

- **OSPF**

- Using OSPF with Multiple Default Gateways
- Using OSPF with BGP
- Several Methods and Implementations

What is BGP?

- **Border Gateway Protocol**
 - Exterior Gateway Protocol
 - Used as the core routing protocol on the Internet
 - Used to interconnect separate diverse networks
 - Uses AS numbers to identify networks
 - AS numbers advertise Prefixes – IP addresses
 - Only way to Properly Multi-Home your Network

What is BGP?

- **Border Gateway Protocol**
 - **AS – Autonomous System**
 - *Collection of Prefixes under the control of one or more operators.*
 - *Typically a single 'network'*
 - *Many prefixes can be in one AS*
 - **Prefixes**
 - *38.114.84.0/22*
 - *12.0.0.0/8*

What is BGP?

- **Border Gateway Protocol**

- **Multihoming**

- *A process to increase the reliability of the internet connection for prefixes*
 - This allows you to connect your network at several points to either the same or multiple upstream providers.
 - Primary way to allow proper in/out operation of public IP space.

- **Advertisements**

- *Your AS will advertise your Prefixes*
 - *The advertisements are sent out to your BGP peers*
 - *These peers will then update their peers – so on and so forth*
 - *You can advertise your IP addresses out multiple internet connections (as long as they are running BGP)*

Why do you need BGP?

- **Multihoming**

- This is the primary method to prevent a single internet upstream connection from going down and still having full public IP connectivity
 - *Advertising the same prefixes out multiple providers provides failover between those providers*
 - *Inbound traffic to those public IPs can go through any of your upstream providers (again as long as they have a BGP connection)*

- **Failover**

- This provides redundancy in your network operations
- A single provider that goes down, does not affect your network operations

What do you need to run BGP?

- AS Number
- IP Prefixes to Advertise
- Upstream – BGP Peer

What do you need to run BGP?

- **AS Number and IP Prefixes**
 - **RIR – Regional Internet Registry's**
 - *ARIN – US/Canada/Parts of the Caribbean*
 - *RIPE NCC – Europe, Middle East, Central Asia*
 - *AfriNIC - - Africa*
 - *APNIC – Asia, Australia, other neighboring countries*
 - *LACNIC – Latin America and parts of the Caribbean*
 - *You DO NOT have to have your own Prefixes in some cases*
- **Upstream – BGP Peer**
 - **Your upstream providers must have the capabilities to preform BGP peering**
 - *Most DSL / Cable / Broadband operators in USA do NOT allow BGP peers on these types of circuits*
 - *MetroE, Fiber Services, >= T1 Circuits will typically provide it*
 - *This is not in all cases, contact your provider to find out*

What are the costs to run BGP?

- **AS Number and IP Prefixes**
 - *AS Number – though ARIN*
 - \$500 to register each AS number
 - \$100 Maintenance fee for each AS per subsequent years

What are the costs to run BGP?

- **AS Number and IP Prefixes**

- *IP Prefixes*

- Blocks of IP addresses assigned to YOUR organization
 - Smallest block being assigned is a /22
 - Blocks Smaller than /20 are \$1,250 per year
 - These IPs will require you to advertise these Prefix's via BGP
 - They can be used with ANY internet provider that offers BGP connectivity
 - You change upstream internet providers, no IP renumbering is required.

What are the costs to run BGP?

- **AS Number and IP Prefixes**

- *IP Prefixes – IPv4 – with ARIN*
 - You will have to prove that you NEED these IP Addresses
 - Provide subnet list for POP locations
 - Provide IP / POP / Customer Information
 - Existing Subnets
 - Must be SWIPed to your business name
 - May need LOA from each provider showing you are assigned those subnets.
 - May need LOA between each provider allowing any existing subnets to be advertised by you.
 - Show Planned usage for the next 3 months only!
 - You must show that you have a plan to use the IP addresses assigned within 3 months .
 - You typically will also have to prove that you are or will be multi-homed

What are the costs to run BGP?

- **So.... With ARIN in the US**
 - **AS number**
 - *\$500 First year*
 - *\$100 each year afterword's*
 - **/22 – 4 /24 blocks – 1020 IP Addresses**
 - *\$1,250 per year*
- **Rates and Fees may vary from your RIR.**

Ready to Run!


- **Your Providers**
 - **Typically will give you a BGP worksheet to fill out**
 - *What is your AS*
 - *What other AS's do you connect to*
 - *What Prefixes are you wishing to advertise*
 - Many Providers filter these!
 - **Return with BGP information**
 - *There AS*
 - *There BGP peer IP*
 - *MD5 Key*
-

BGP Prefix Feeds / Connections

- **Default Only**
 - Provider sends only a default route via BGP
- **Summary**
 - Provides only a summarization of global routing table.
 - Typically a few hundred routes.
 - Depends on the provider, some summarize up to /8s some summarize down to /16s

BGP Prefix Feeds / Connections

- **Full Table**
 - Your provider will send you the entire Global Routing Table



How BGP Selects Paths

- **AS Numbers ONLY!**
 - An AS can be
 - *5 Routers*
 - *200 Routers*
 - An AS can be
 - *100 routers in a single Town*
 - *20 routers across the mid-west*
- **BGP Only looks at AS Hop Count!**
 - The number of routers you go through does NOT matter!
 - Typically you are within 5 AS hops from anywhere on the internet
 - A Global internet provider will be typically closer than a smaller internet provider


BGP Implementations

- **Things to Know**
 - **Some providers will allow you to announce their prefix's that are assigned to you**
 - *If you disconnect that provider, they typically want those prefixes back!*
 - **Most providers require you to tell them, what prefix's you are going to announce.**
 - *They filter based on the prefix's you tell them, so if you want to add a new prefix, you may have to tell your peers prior to announcing.*
 - *You may need to have a LOA if you are announcing another provider's prefix.*

BGP Implementation Examples

BGP Implementations

- **Single Router – Two Upstream Providers**
 - Very common in Datacenters
 - Two Providers Equal Bandwidth
 - Each offer BGP connections



BGP Implementations

- **Single Router – Two Upstream Providers**
 - **Suggestions**
 - *Full Routing Tables from each provide*
 - Allows us to identify and use the most cost effective route from each provider
 - *Advertisements*
 - Since both peers are the same speed, we typically don't need to change our inbound routing.

BGP Implementations

- **Single Router – Two Upstream Providers**
 - **BGP Configuration**
 - *Do we require Default routes?*
 - All Prefixes on the Global Routing table are learned. No other routes should be needed!
 - *Dual Global Routing Tables*
 - BGP will use the AS hop count, some prefixes will go out though ISP1 and some will go through ISP2 based simply on AS Hop Count
 - Inbound Traffic will be the same

BGP Implementations

- Two Routers -- Two Peers – Same Physical Location




BGP Implementations

- **Two Routers -- Two Peers – Same Physical Location**
 - **Suggestions**
 - *Full Routes still needed*
 - *Create iBGP peer between two routers to share information*
 - This is done by creating a BGP peer with same AS numbers on both ends
 - *Extra router hop does NOT matter*
 - *If best AS path is through Router B, and packet comes in Router A, Router A will send packet to Router B*
 - Typically you will want a dedicated Ethernet or virtual connection between the two routers

BGP Implementations

- Two Routers -- Two Peers – Same Physical Location



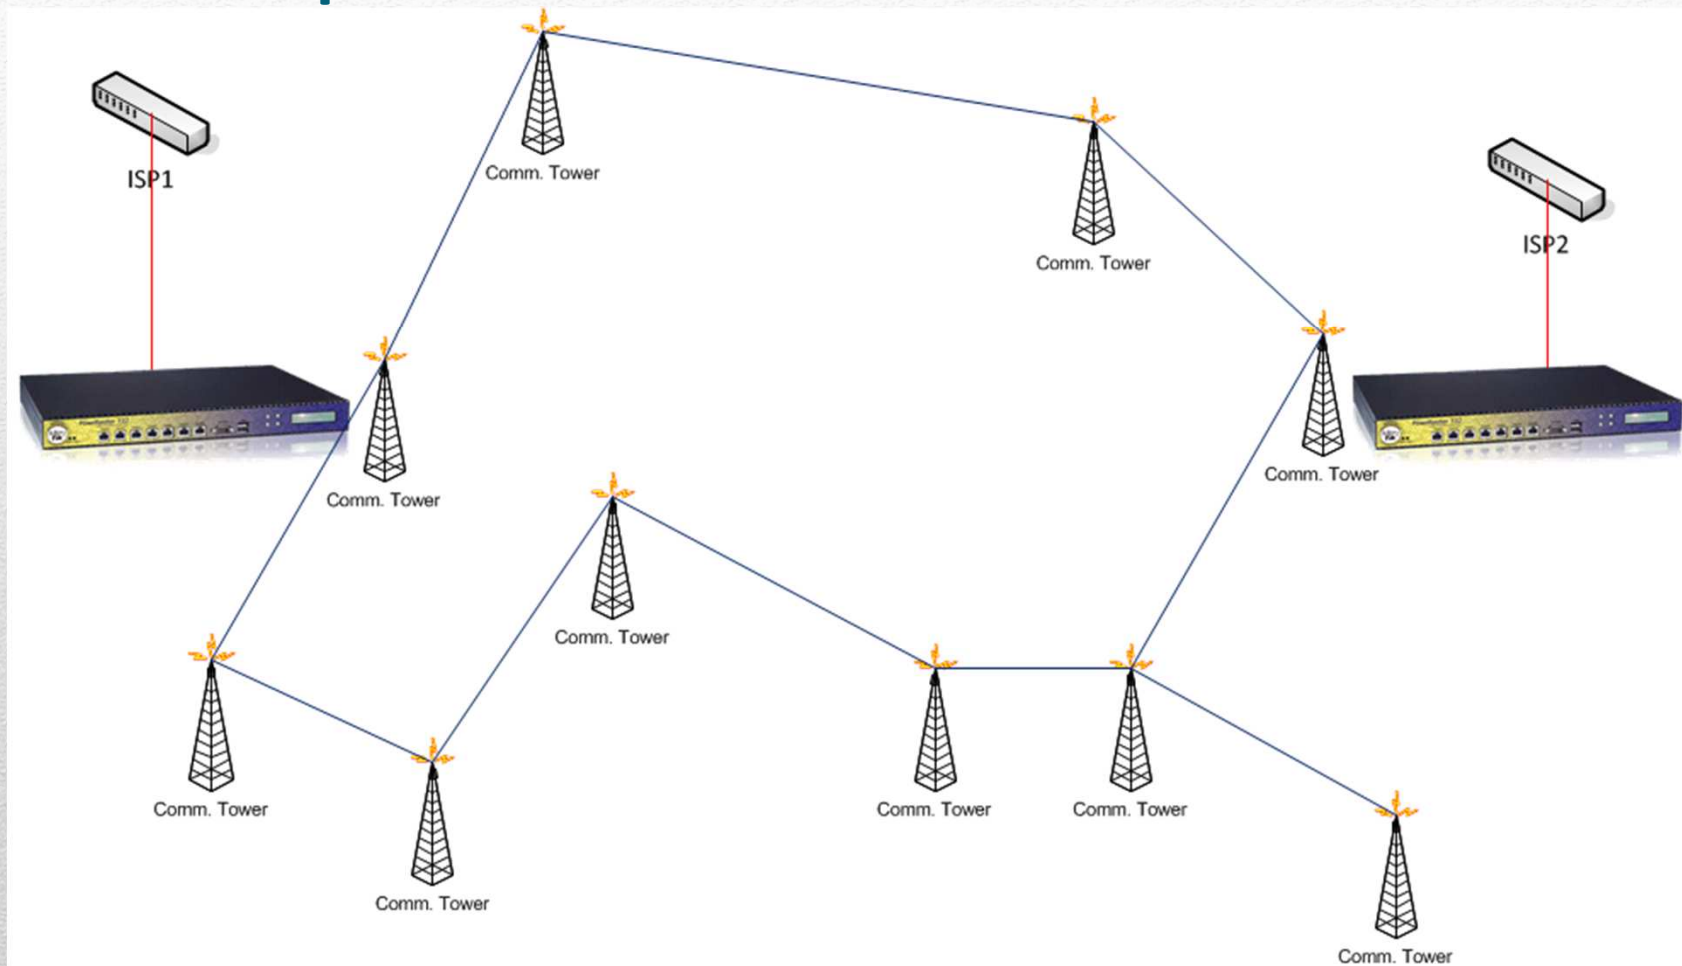
BGP

- **Use BGP to select best inbound/outbound route from your network to others**
 - Remember BGP is a EGP (Exterior Gateway Protocol)
 - *Designed to interconnect NETWORKS*
 - **Use OSPF inside your network**
 - To get traffic to and from your network edges
-

OSPF Implementation Examples

OSPF Implementation

- Multiple Diverse Internet Connections

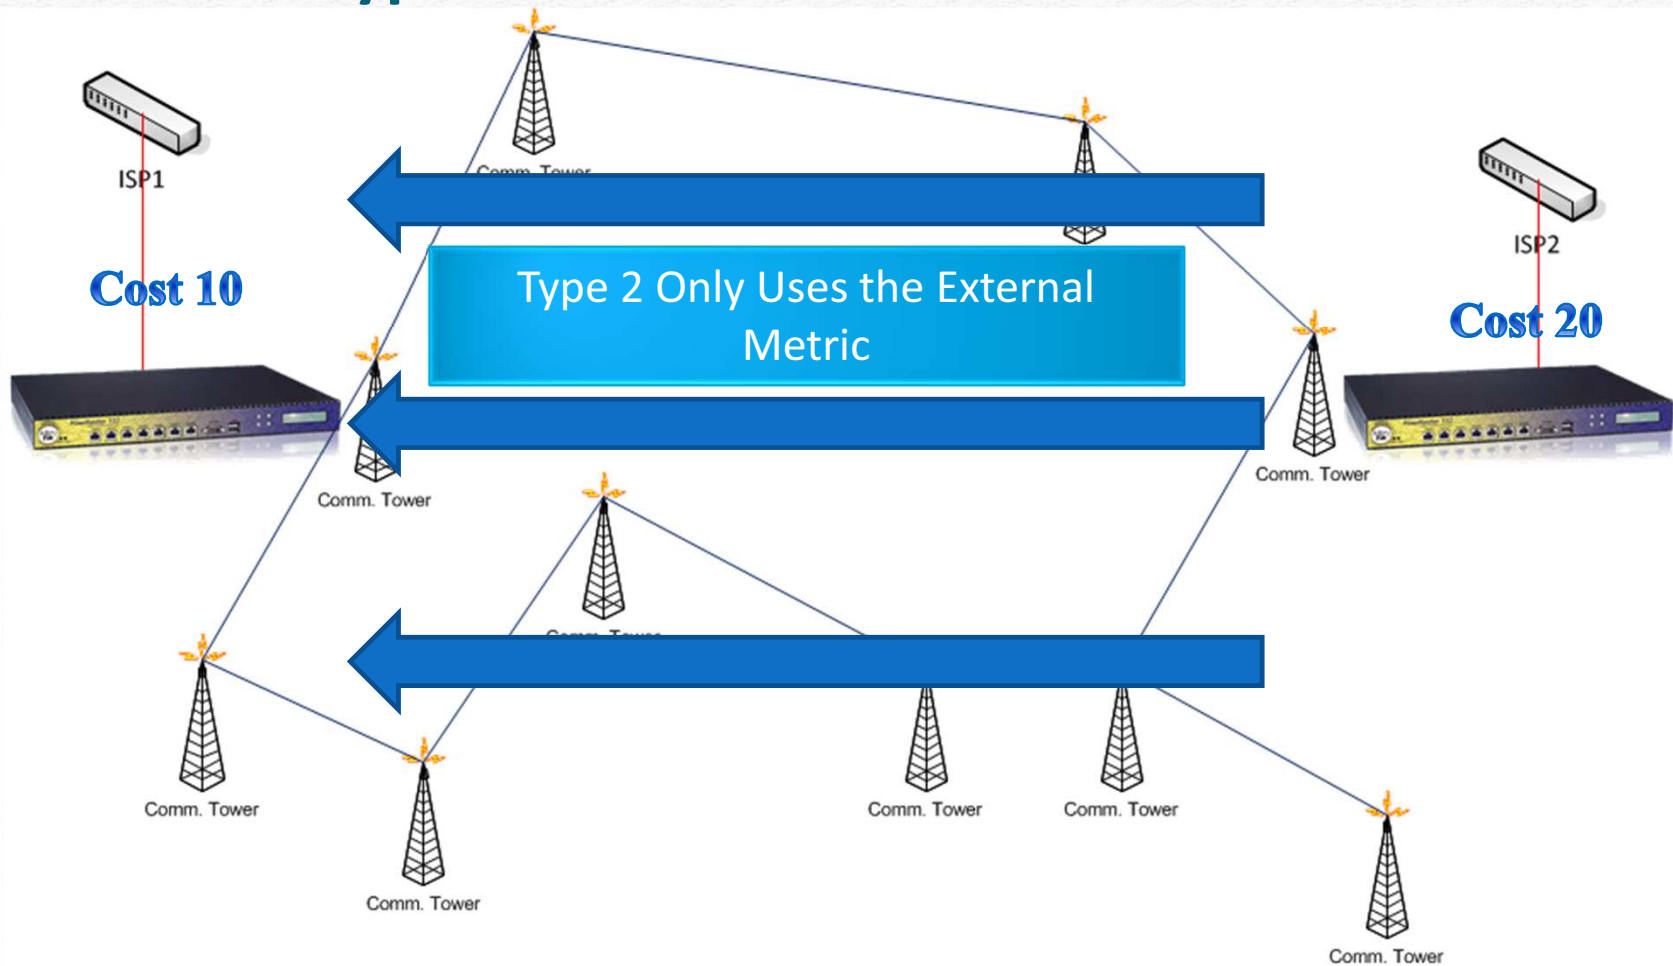


OSPF Implementation

- **Multiple Diverse Internet Connections**
 - **OSPF Type 1**
 - *Uses both Internal and External Metrics*
 - **OSPF Type 2**
 - *Uses ONLY External Metrics*
- **Suggestions**
 - *Set Default route to OSPF Type 1*
 - This way you can SPLIT your network
 - Some traffic though ISP1 and some though ISP2
 - You can control this by increasing your link costs across interfaces
 - *Set Default route to OSPF Type 2*
 - Only one Default gateway will be used, unless unavailable or connectivity failure.


BGP Implementations

- OSPF Type 2 Default Route



OSPF Implementation

- OSPF Type 1 Default Route



OSPF Implementation

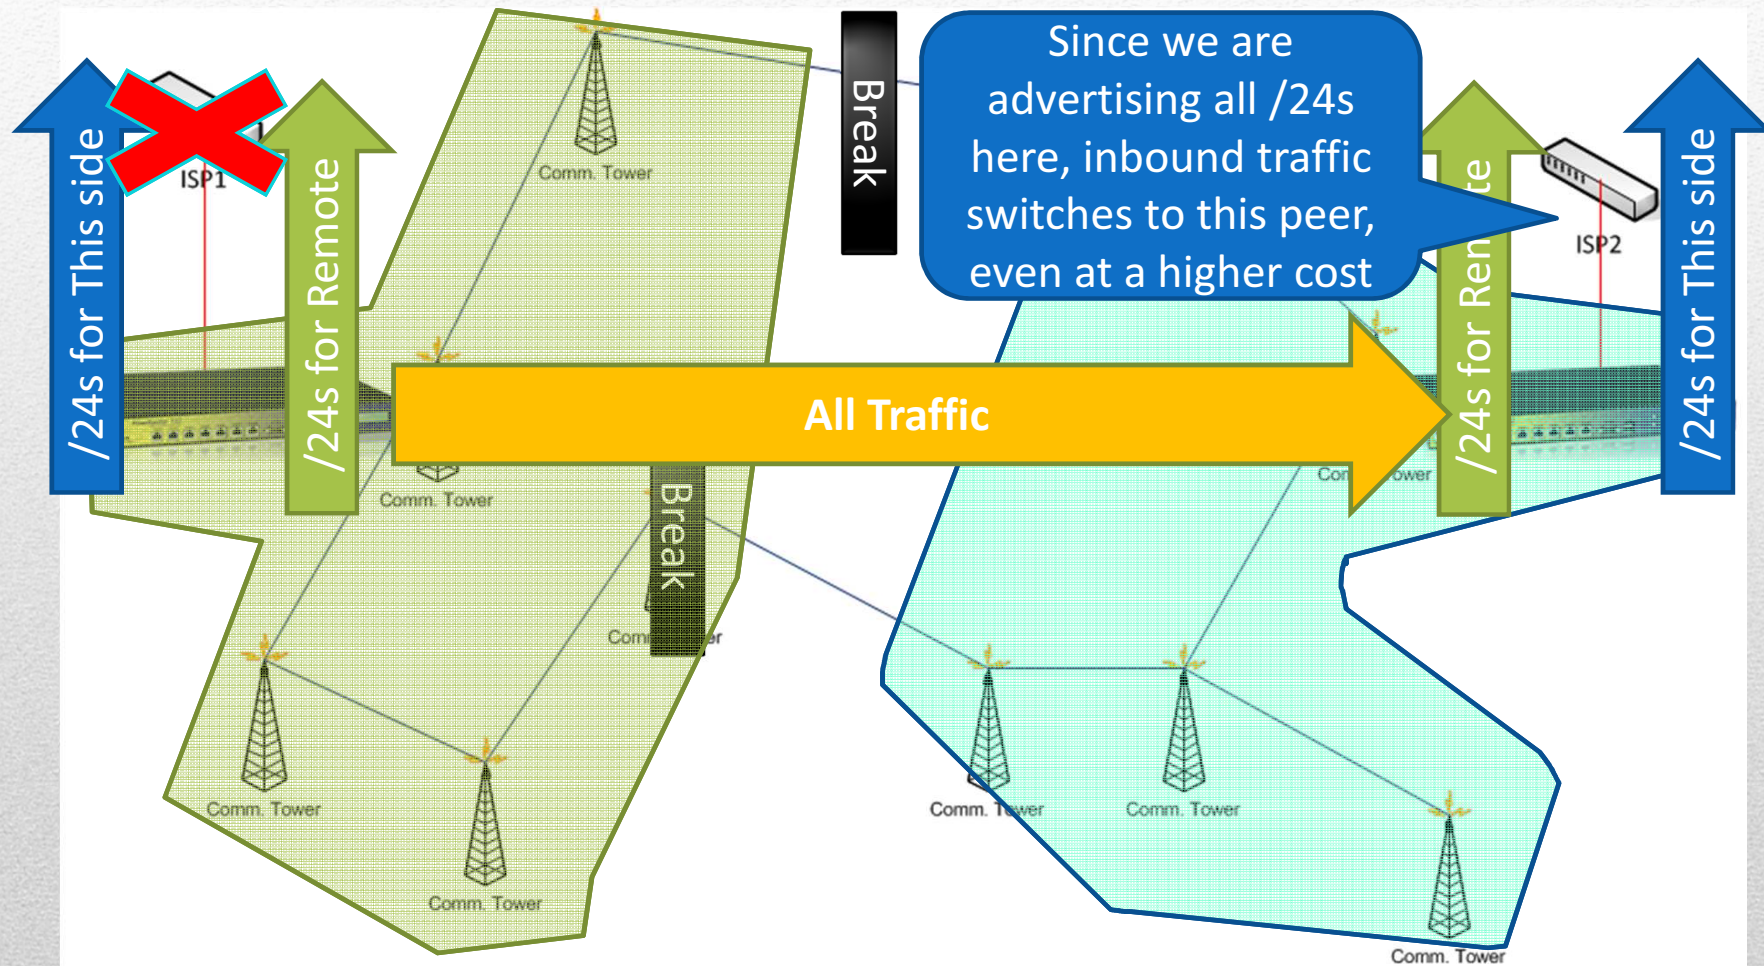
- **OSPF Type 1 Default Route**
 - **Breaks Network**
 - *In the event of failure on one gateway, all traffic goes to the other.*
 - *Also Useful for other BGP Routing*
 - **By using Two default routes you can balance out your outbound network traffic between both gateways**
 - *Does not balance inbound traffic*
 - *In many cases there would be NATs at each end*

BGP/OSPF Implementation

- **How do we use the OSPF Type 1 Default routes with BGP inbound routing?**
 - **Suggestions**
 - *Separate your prefixes to each “section” of your physical network topology*
 - *Only need Default routes from peers*
 - *The Smallest prefix typically allowed to advertise is /24*
 - Do not put subnets of /24s on different segments of your network
 - Always Keep /24s to either section
 - *Advertise /24 prefixes out their respective BGP peer*
 - *Advertise the other peers /24s with Prepends*
 - Adds AS cost to advertisement
 - Typically 3-5 prepends are sufficient

BGP/OSPF Implementation

- Two Diverse BGP Peers – Two Providers

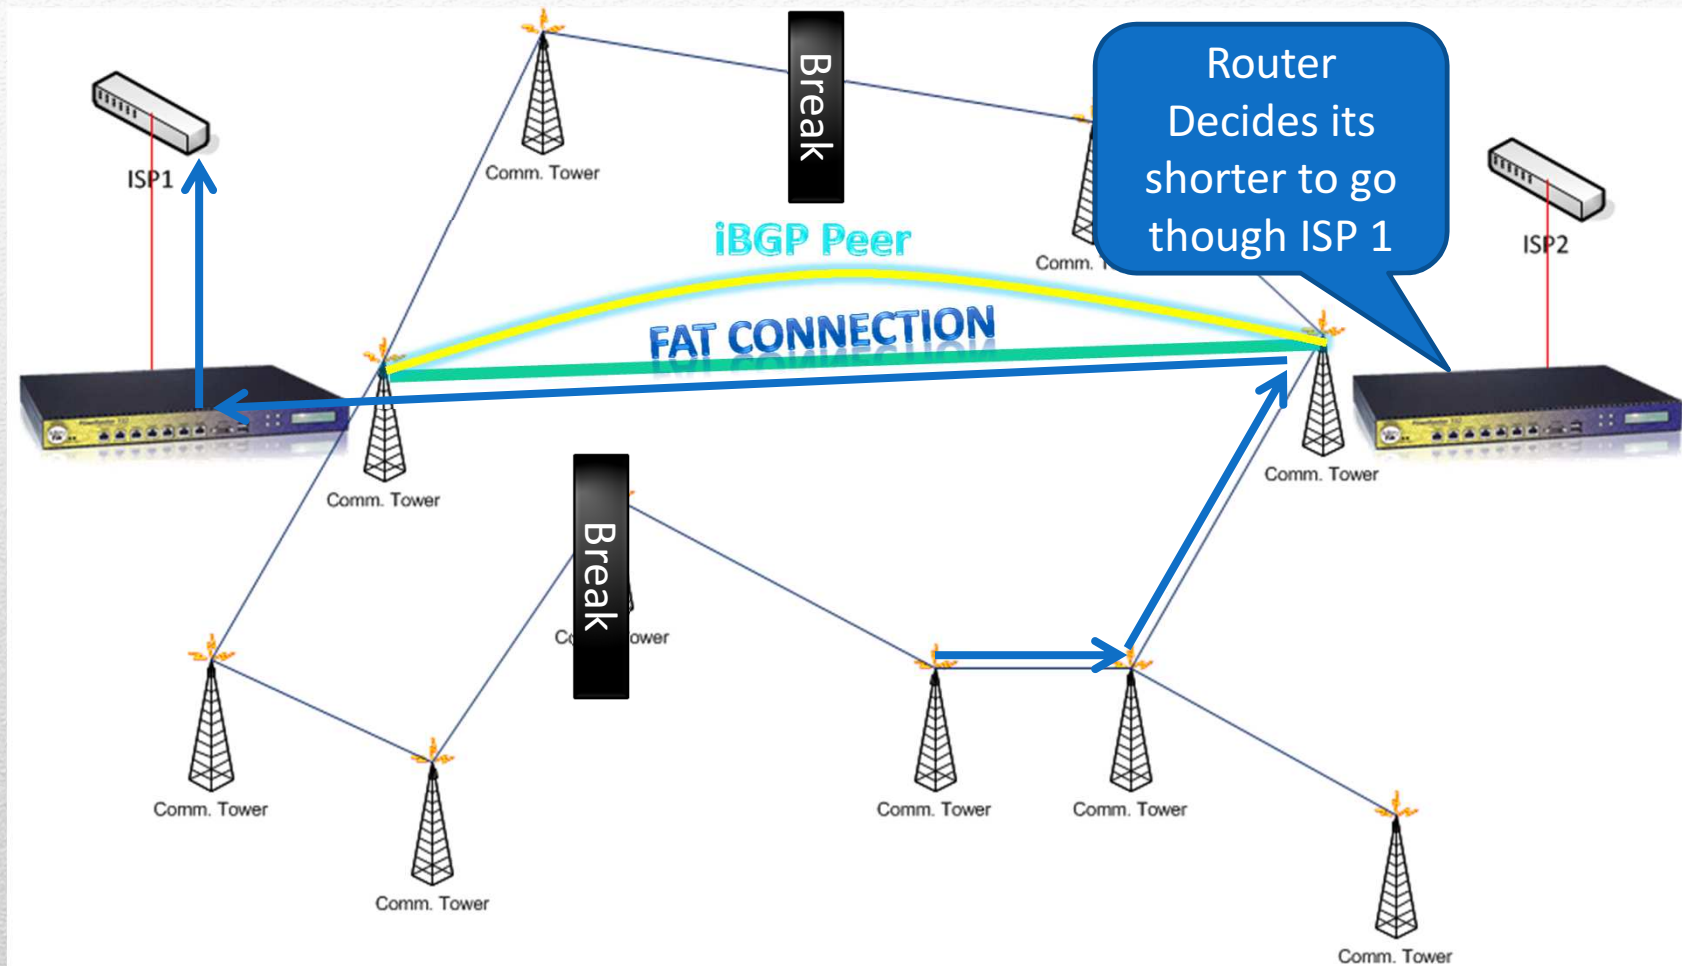


BGP/OSPF Implementation

- **Two Diverse BGP Peers – Two Providers**
 - **What if we add a “FAT PIPE” Between our diverse Connections?**
 - *We can now utilize full routing tables at each BGP peer*
 - *Routing occurs just like a single router with multiple peers*
 - Extra router hops within the same AS does not matter
 - *Best Path will be selected based on destination AS count*

BGP/OSPF Implementation

- Two Diverse BGP Peer – Two Providers



Other Information

- **BGP**

- Complete books entirely on traffic engineering, using things like communities, prepends, changing costs, etc.
 - *Recommend that you setup lab environments to test prior to deployment*
 - *Consult with MikroTik Certified Consultant with BGP Experience*

- **OSPF**

- Books on OSPF deployments, areas, route summarization, path costing, full duplexing over half-duplex links.
 - *Refer to these*

Your Presenter

- Dennis Burgess

Link Technologies, Inc.

- Contact Information
 - <http://www.linktechs.net>
 - support@linktechs.net
 - Office: 314-735-0270
- “Learn RouterOS” Book
 - www.routerosbook.com