EFFICIENT AND SECURE INTERNET ROUTING: BEST PRACTICES AND AUTOMATION

Presented by Rafeeun Noby Babir

INTRODUCTION

Topic: Efficient and Secure Internet Routing: Best Practices and Automation

Agenda:

Overview of Internet routing,

Routing Best Current Practices (BCP),

Automation in large-scale network deployment

WHAT IS INTERNET ROUTING?

- Internet routing is the process of forwarding data packets from one network to another through a network of routers.
- The Internet routing system
  - **Autonomous systems (AS)** that exchange routing information
  - Standard protocols such as **Border Gateway Protocol (BGP)**

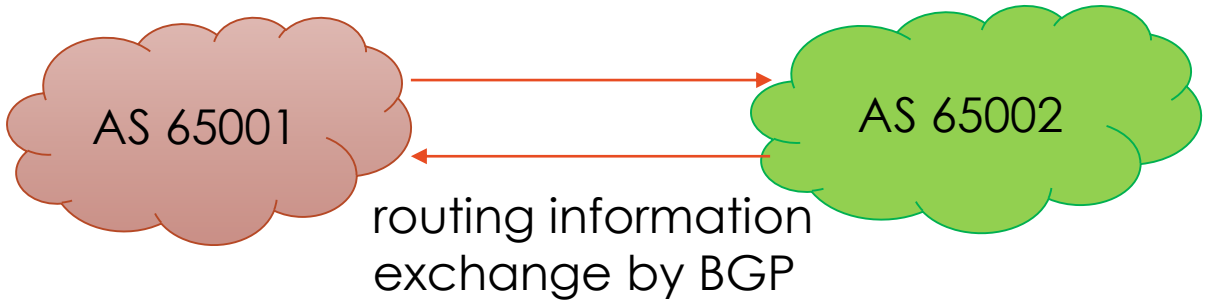

CHALLENGES WITH INTERNET ROUTING

- Internet routing is a **complex process**
  - multiple entities
  - number of challenges that need to be addressed
- Internet routing is the **lack of**
  - visibility
  - control over the routing paths
- Risk of **security threats**
  - routing attacks
  - prefix hijacking
  - route leaks

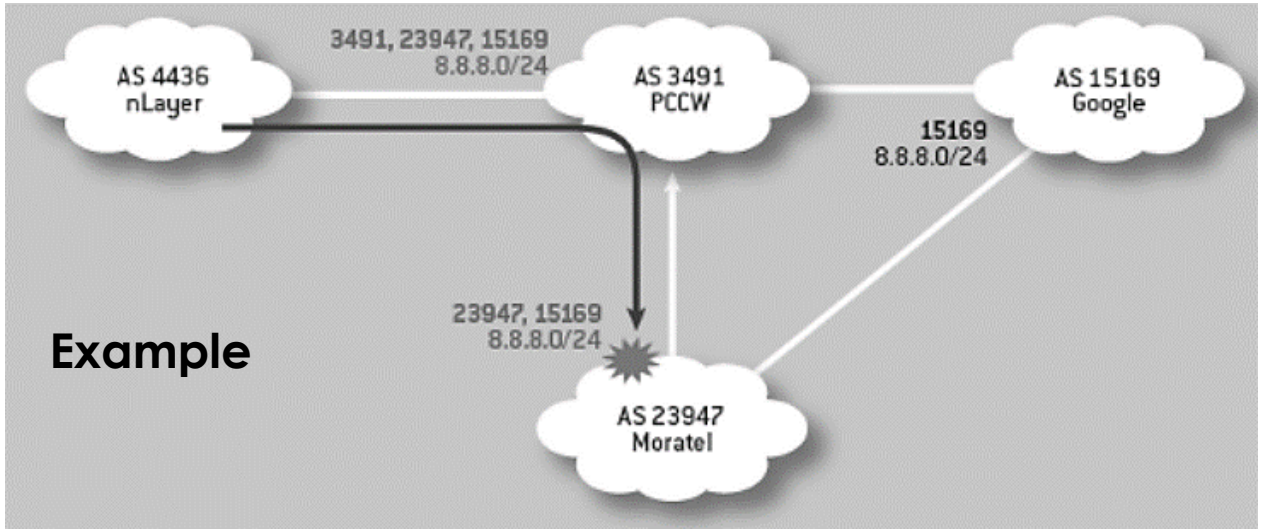

On November 6, 2012, a misconfiguration at Moratel did just that, "leaking" the route of 8.8.8.0/24

ROUTING BEST CURRENT PRACTICES (BCP)

- **Routing Best Current Practices (BCP) are:**
  - Review of recommendations to ensure the BGP is Configured optimally, Operating optimally, Configured securely, Operating securely
  - Covers routing security, route filtering, and operational practices
  - Designed to help network operators improve reliability and security of routing infrastructure
- **Examples of BCPs include**:
  - Route filtering to prevent route leaks
  - Prefix filtering to prevent prefix hijacking
  - Secure BGP to prevent attacks

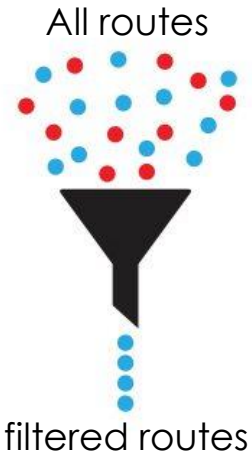

INTERNAL ROUTING PROTOCOL (IGP) AND BGP USAGE

- IGPs (OSPF/IS-IS):
  - used for carrying infrastructure addresses, not Internet or customer prefixes
  - Design goal: minimize the number of prefixes in IGP for scalability and rapid convergence
- BGP
  - used both internally (iBGP) and externally (eBGP)
  - iBGP used to carry some/all Internet prefixes across the backbone and customer prefixes
  - eBGP used to exchange prefixes with other ASes and implement routing policy
- DO NOT
  - distribute BGP prefixes into an IGP or distribute IGP routes into BGP
  - use an IGP to carry customer prefixes

INTERNAL ROUTING PROTOCOL (IGP) AND BGP USAGE

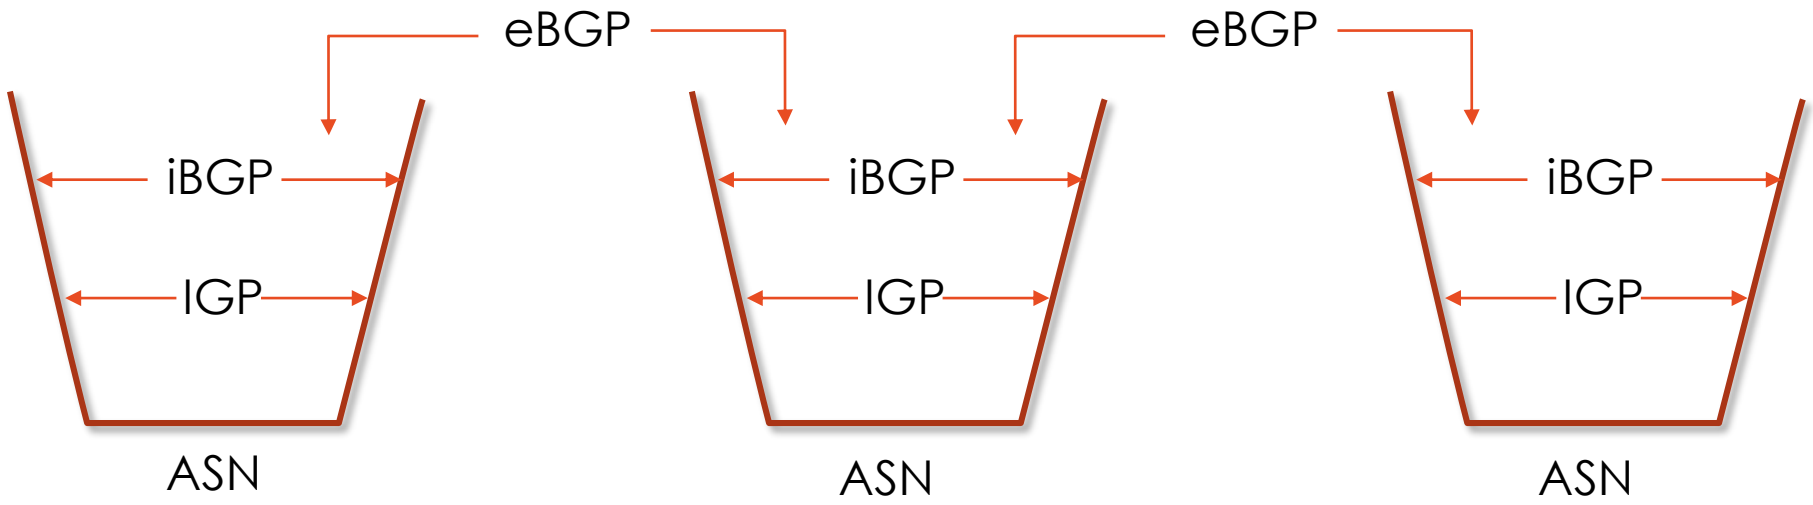

PREFIX INBOUND FILTER

- When receiving prefixes from other ASNs, there are three scenarios to consider:
  - Customer talking BGP
  - Peer talking BGP
  - Upstream/Transit talking BGP
- Each of these scenarios has different filtering requirements and should be considered separately.
- Effective filtering is important to prevent the propagation of bad or malicious routes across the Internet.
- **Do Not**
  - accept ipv4 prefixes longer than /24
  - accept ipv6 prefixes longer than /48

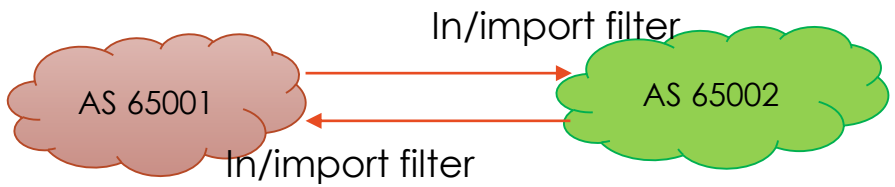

PREFIX INBOUND FILTER: FROM CUSTOMER

- ISPs should only accept prefixes that have been assigned or allocated to their downstream customers.
- If an ISP has assigned address space to its customer, the customer is entitled to announce it back to the ISP.
- If the ISP has NOT assigned address space to its customer, then:
  - Check in the five RIR databases (ie RADB) to see if the address space has really been assigned to the customer.
- Validate prefix announcements received from EBGP peers
  - Drop *invalid*, accept *valid*, low priority for those with no ROA

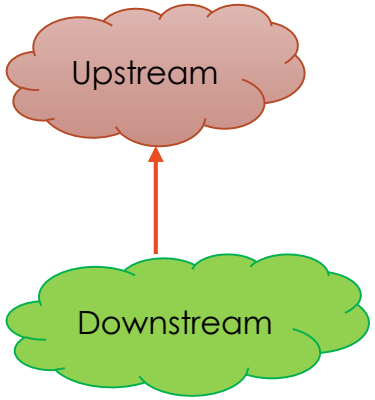

PREFIX INBOUND FILTER: FROM CUSTOMER

- An example of RADB output

```
route:          103.15.246.0/24
descr:          Route Object for 103.15.246.0/24
origin:         AS58717
country:        BD
mnt-lower:      MAINT-SUMMITCOMMUNICATIONS-BD
mnt-routes:     MAINT-SUMMITCOMMUNICATIONS-BD
mnt-by:         MAINT-SUMMITCOMMUNICATIONS-BD
changed:        shahidullah.kaisar@summitcommunications.net 20140303
source:         APNIC
```

This prefix is allowed to originated from that ASN

- RPKI server Configuration example for Cisco:

```
router bgp 65002
 bgp log-neighbor-changes
 bgp rpki server tcp 100.68.3.6 port 3323 refresh 900
```

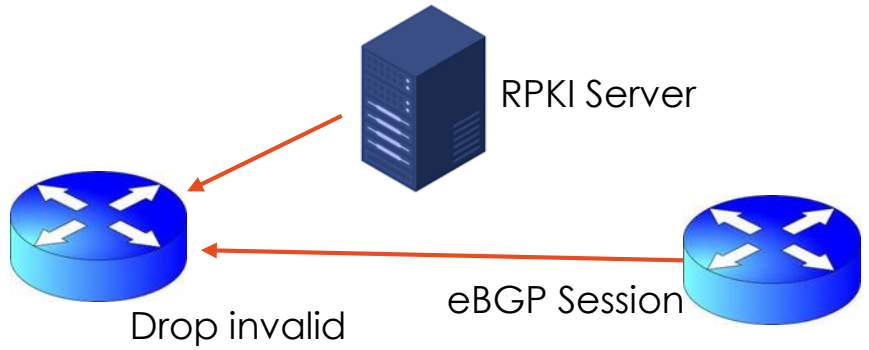

PREFIX INBOUND FILTER: FROM CUSTOMER

- An example of customer configuration in upstream (AS65001):

```
router bgp 65001
neighbor 100.68.1.2 remote-as 65002
 !
 address-family ipv4
  neighbor 100.68.1.2 activate
  neighbor 100.68.1.2 prefix-list customer in
  neighbor 100.68.1.2 prefix-list default out
 exit-address-family
!
ip prefix-list customer seq 5 permit 100.68.2.0/23 le 24
!
ip prefix-list default seq 5 permit 0.0.0.0/0
```

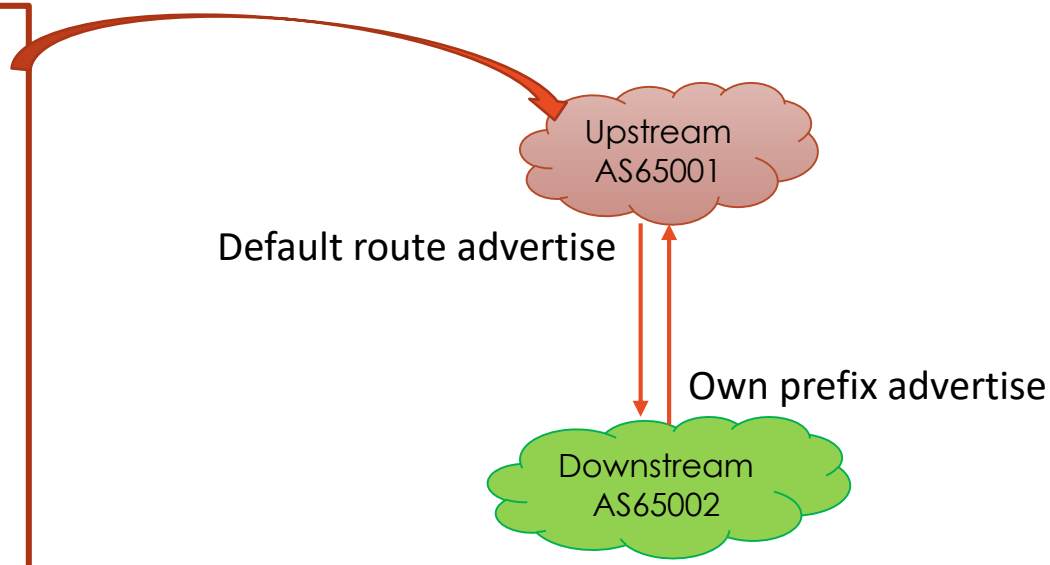

PREFIX INBOUND FILTER: FROM UPSTREAM/PEER

- Upstream/Transit Provider is an ISP you pay for transit to the whole Internet
  - Receiving prefixes from them is not recommended unless necessary for traffic engineering
  - Simplify routing by asking them to either originate a default route or announce one prefix you can use as default.

- Peers are ISPs that exchange prefixes with each other
  - You only accept and announce the prefixes that you have agreed upon with your peer

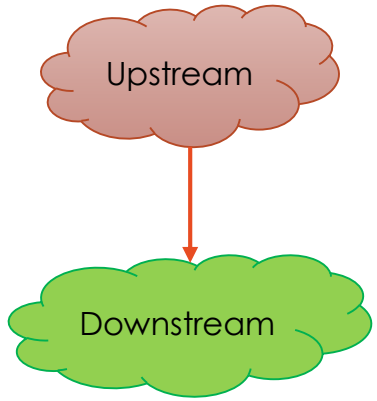

PREFIX INBOUND FILTER: FROM UPSTREAM/PEER

- An example of upstream configuration in customer router:

```
router bgp 65002
neighbor 100.68.1.1 remote-as 65001
 !
 address-family ipv4
  neighbor 100.68.1.1 activate
  neighbor 100.68.1.1 prefix-list upstream in
  neighbor 100.68.1.1 prefix-list export out
 exit-address-family
!
ip prefix-list export seq 5 permit 100.68.2.0/24
ip prefix-list export seq 10 permit 100.68.3.0/24
!
ip prefix-list upstream seq 5 permit 0.0.0.0/0
```

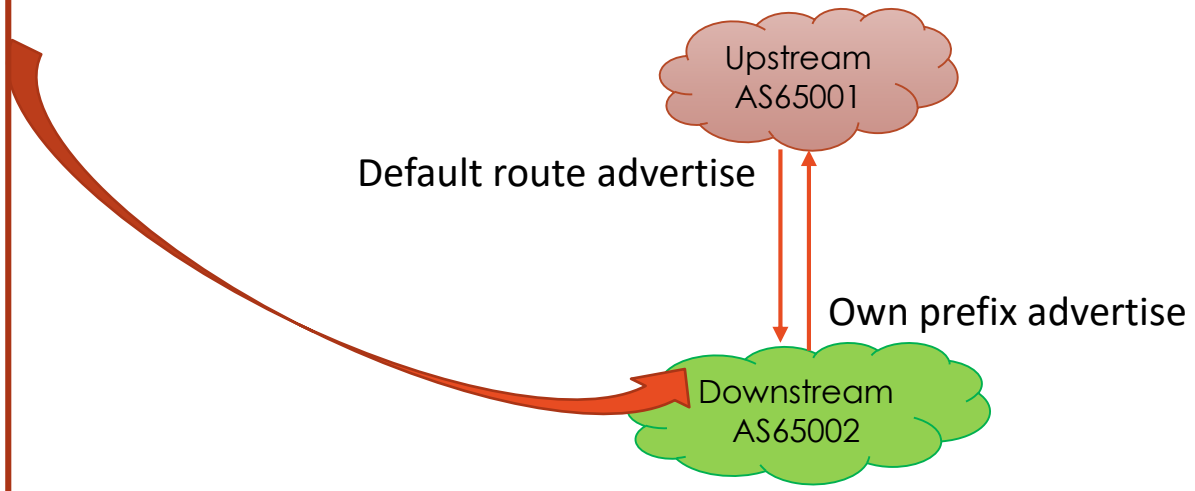

PREFIX INBOUND FILTER: BOGON

- Bogon prefix is a route that
  - should never appear in the Internet routing table
  - often used as the source addresses of DDoS attacks
- Filter or deny bogon prefixes from both upstream and customer routes
- Team Cymru Bogon list for reference:
  - https://www.team-cymru.com/bogon-networks

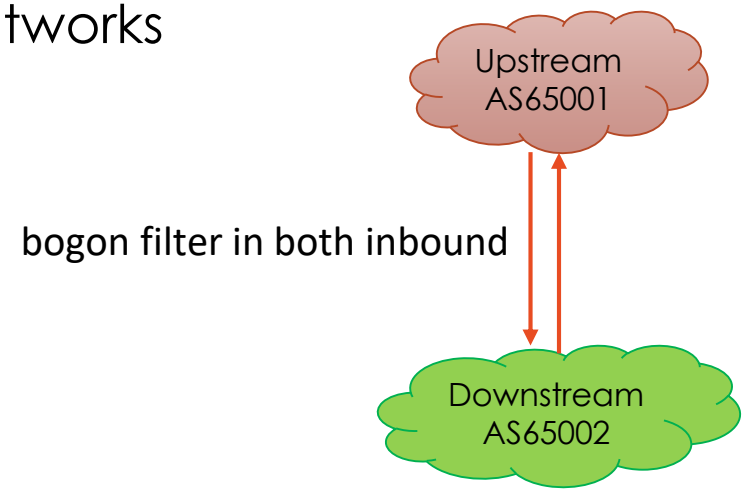

BGP MAXIMUM PREFIX LIMIT

- Configure the maximum number of prefixes a BGP router will receive from a peer
  - Set the "maximum prefixes" to be 2xN, where N is the expected number of prefixes
  - If the limit is exceeded, the router will warn or tear down the BGP session
  - Even if receiving the full BGP table, setting a limit is still a good idea
  - This prevents against major accidental leaks.

- Configuration example for Cisco:

```
router bgp [AS_number]
 address-family ipv4
 neighbor [IP_address] maximum-prefix [number of prefixes] warning-only/restart [restart interval]
```

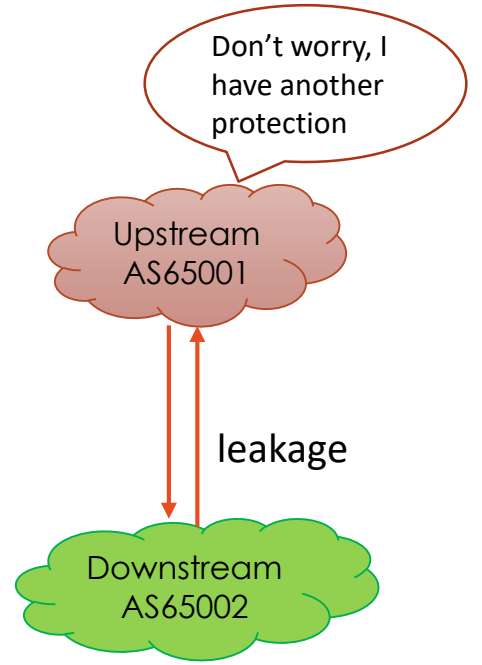

AUTOMATION IN LARGE-SCALE NETWORK DEPLOYMENT

- Automation is essential for managing large-scale network deployments in the growing Internet landscape.
  - Helps reduce the risk of errors and increases speed and efficiency.

- There are various network automation tools and platforms available:
  - Paramiko,
  - Netmiko,
  - NAPALM,
  - **Ansible**
  - Nornir

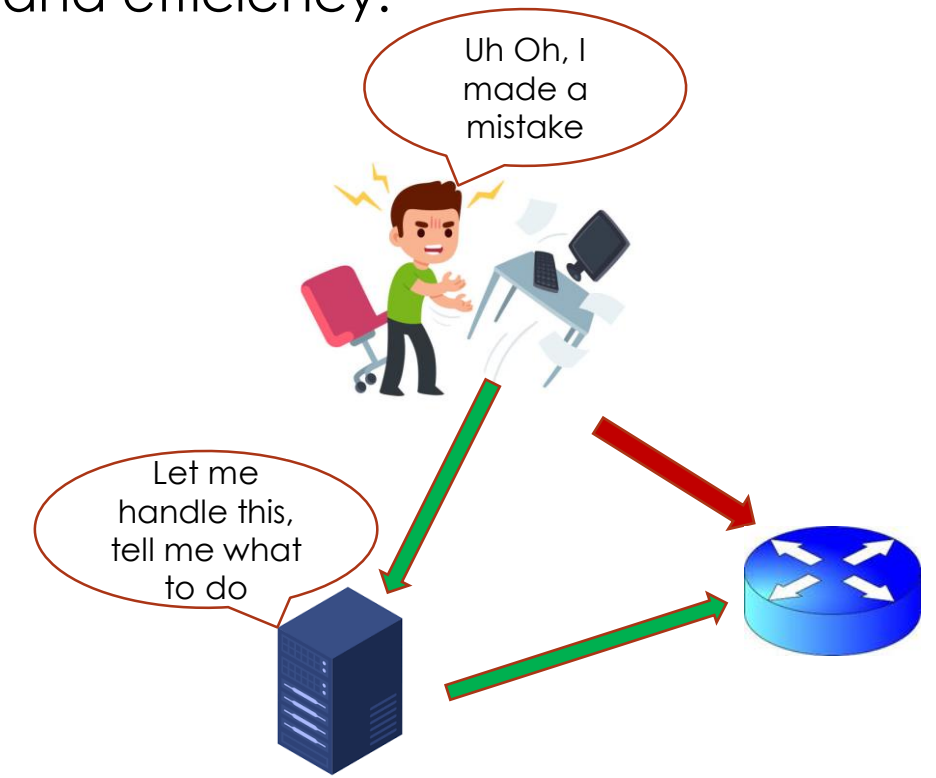

BENEFITS OF AUTOMATION IN NETWORK DEPLOYMENT

- Automation is essential for managing large-scale network deployments in the growing Internet landscape with complex routing systems.
- Automation benefits for network operators:

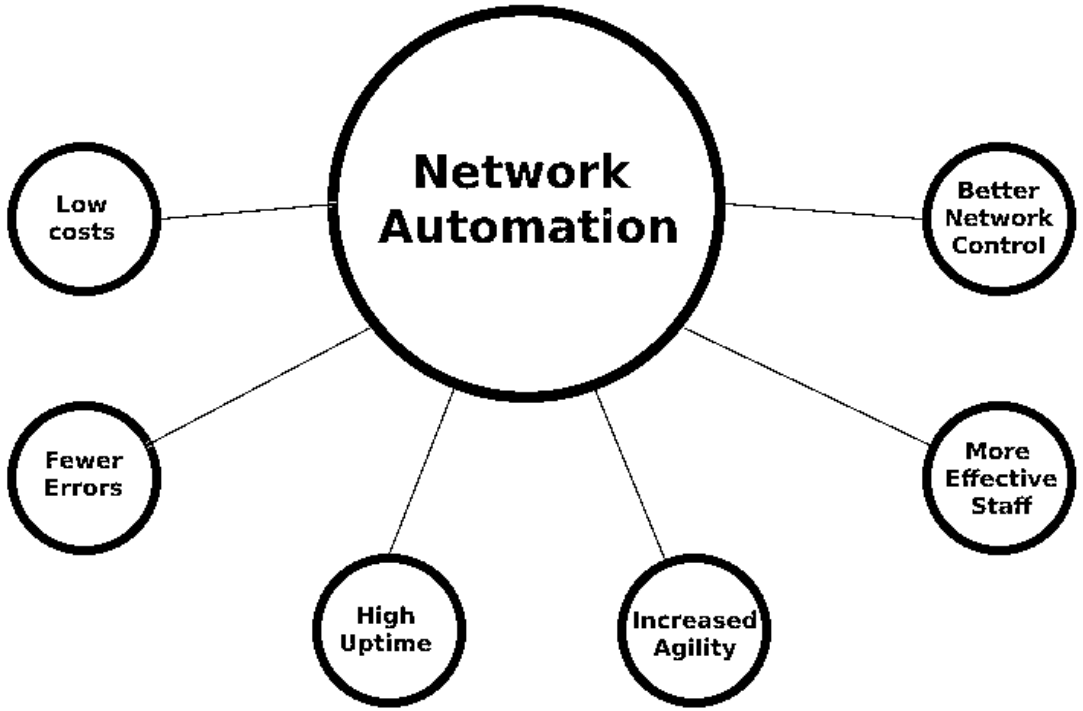

SUMMARY

- Internet routing is a complex process.
- Requires careful management and best practices for reliable and secure connectivity.
- Best practices for network operators:
  - Follow Routing Best Current Practices (BCP).
  - Leverage automation tools.
- Benefits of following best practices:
  - Improved efficiency, reliability, and security of routing infrastructure.