Scaling Broadband Networks

Cameron Dry
iiNet

Introduction

- Network scalability in a broadband subscriber environment
- Real world

  Challenges

  Solutions
- Every SP has specific constraints

  ∴ YMMV

Introduction

- Problem ? Where ?
- Solution to scaling networks is easy.....
- Simply throw more resources at it

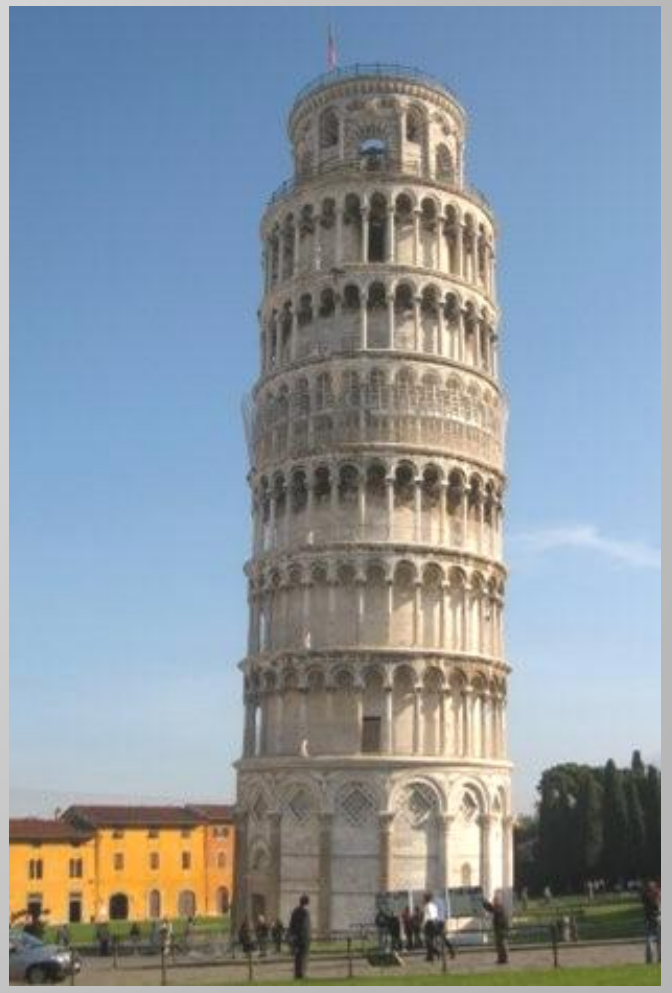

Introduction

- However, must consider…
  - Real-world constraints
  - Factors outside of our control
  - Commercial Reality™

Introduction

- Definition of “scale” in an iiNet context

> 230k BB users

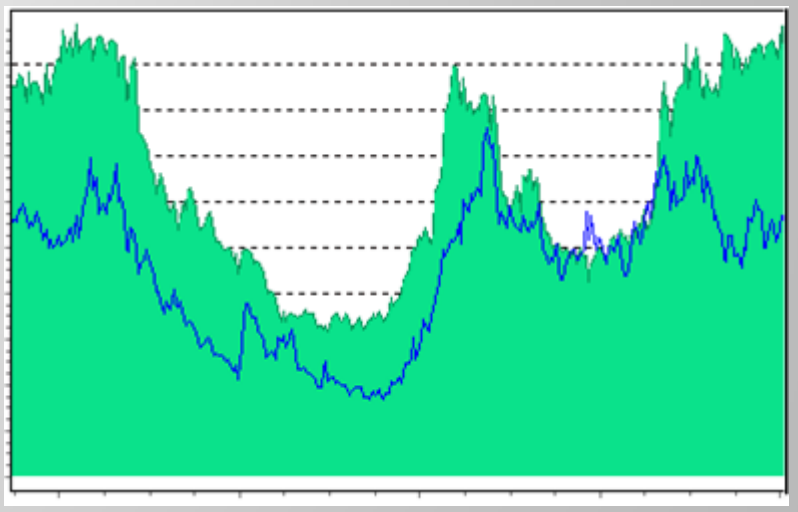

>15 Gbps IP

Agenda

- Scalability in the context of.....
  - **Authentication**
  - Address Allocation
  - Accounting
  - Bandwidth

Authentication

- Requirements and constraints
  - Authenticate all BB users in a State within 5 minutes
  - 75,000 users in a PoP → 250 sessions/sec
  - Maintain RADIUS technology
  - No manual intervention

Authentication

- Several optimisations required to achieve customer authentication goal
  - Reduce belligerent auth requests
  - RADIUS transaction efficiency
  - Extended radius source ports
  - Local address assignment

Authentication

Reduction of belligerent authentication requests

- Around 200K nonsensical access-rejects per day

- Deployment of PPP connection throttling
  - Acts as DoS protection
  - e.g. after 3 failed authentication requests in 60 seconds, BRAS ignores that user for 5 minutes

- Future implementations will eliminate this issue altogether – 2 stage approach

Authentication

- Several optimisations required to achieve customer authentication goal
  - ✓Reduce belligerent auth requests
  - RADIUS transaction efficiency
  - Extended radius source ports
  - Local address assignment

Authentication

Radius transaction efficiency

- Originally there was a significant quantity of per-user information being sent in authorisation transaction
    - QoS policies – many flavours
    - Filters – standard bogon + user-selectable
- Aim to reduce quantity of information sent from RADIUS to BRAS by using locally defined policies.

Authentication

RADIUS

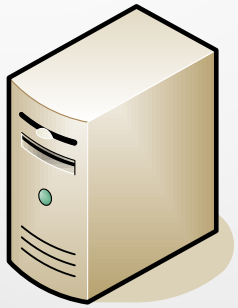

ACL sent for every user

BRAS

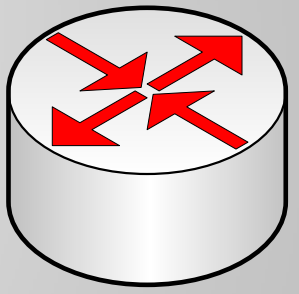

No ACL config on BRAS

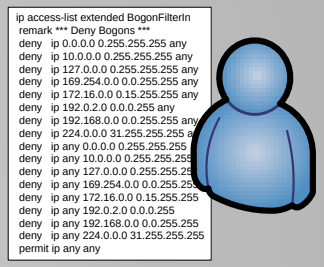

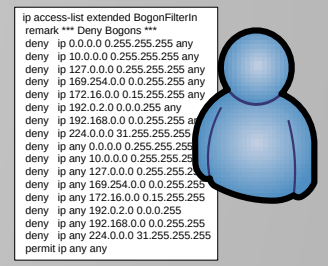

Authentication

RADIUS

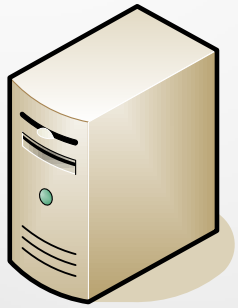

Only filter names are defined on radius

joe user: ip:inacl = “Filter#1”
bob user: ip:inacl = “Filter#2”

BRAS

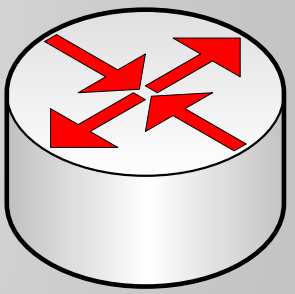

ACLs defined on BRAS

```
ip access-list extended BogonFilterIn
 remark *** Deny Bogons ***
 deny   ip 0.0.0.0 0.255.255.255 any
 deny   ip 10.0.0.0 0.255.255.255 any
 deny   ip 127.0.0.0 0.255.255.255 any
 deny   ip 169.254.0.0 0.0.255.255 any
 deny   ip 172.16.0.0 0.15.255.255 any
 deny   ip 192.0.2.0 0.0.0.255 any
 deny   ip 192.168.0.0 0.0.255.255 any
 deny   ip 224.0.0.0 31.255.255.255 any
 deny   ip any 0.0.0.0 0.255.255.255
 deny   ip any 10.0.0.0 0.255.255.255
 deny   ip any 127.0.0.0 0.255.255.255
 deny   ip any 169.254.0.0 0.0.255.255
 deny   ip any 172.16.0.0 0.15.255.255
 deny   ip any 192.0.2.0 0.0.0.255
 deny   ip any 192.168.0.0 0.0.255.255
 deny   ip any 224.0.0.0 31.255.255.255
 permit ip any any
```

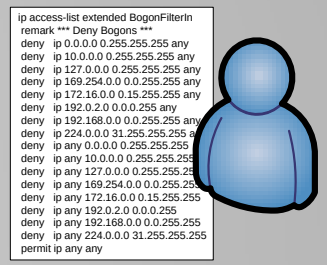

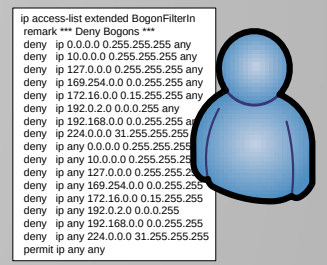

Authentication

- Several optimisations required to achieve customer authentication goal
  - ✓Reduce belligerent auth requests
  - ✓RADIUS transaction efficiency
  - Extended radius source ports
  - Local address assignment

Authentication

Extended radius source ports

- By default, a single source port (1645) is used by the BRAS for all authentication requests
- Max 256 access requests can be outstanding at any one time

```
AUTH: Duplicate authentication request id=12 already in progress
CHAP: I RESPONSE id 12 len 28 from "joeuser"
AUTH: Duplicate authentication request id=12 already in progress
CHAP: I RESPONSE id 12 len 28 from "joeuser"
AUTH: Duplicate authentication request id=12 already in progress
```

Authentication

- Use extended radius source ports 21645-21844
- BRAS cycles sequentially through source ports

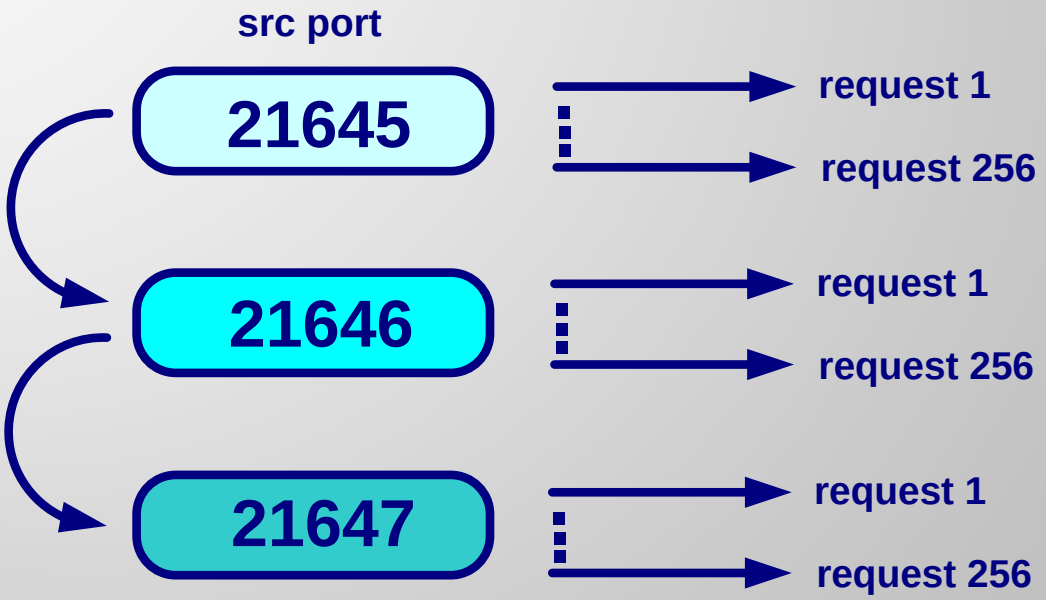

- Now have 200*256 authentication requests in progress

Authentication

- Several optimisations required to achieve customer authentication goal
  - ✓Reduce belligerent auth requests
  - ✓RADIUS transaction efficiency
  - ✓Extended radius source ports
  - Local address assignment

Authentication

Local Address Assignment

- Originally all IP addresses were assigned from a centralised database
  - Simple capacity management
  - Efficient use of address space
  - Operational consistency

Authentication

- Retrieving the IP address from the database was the performance bottleneck

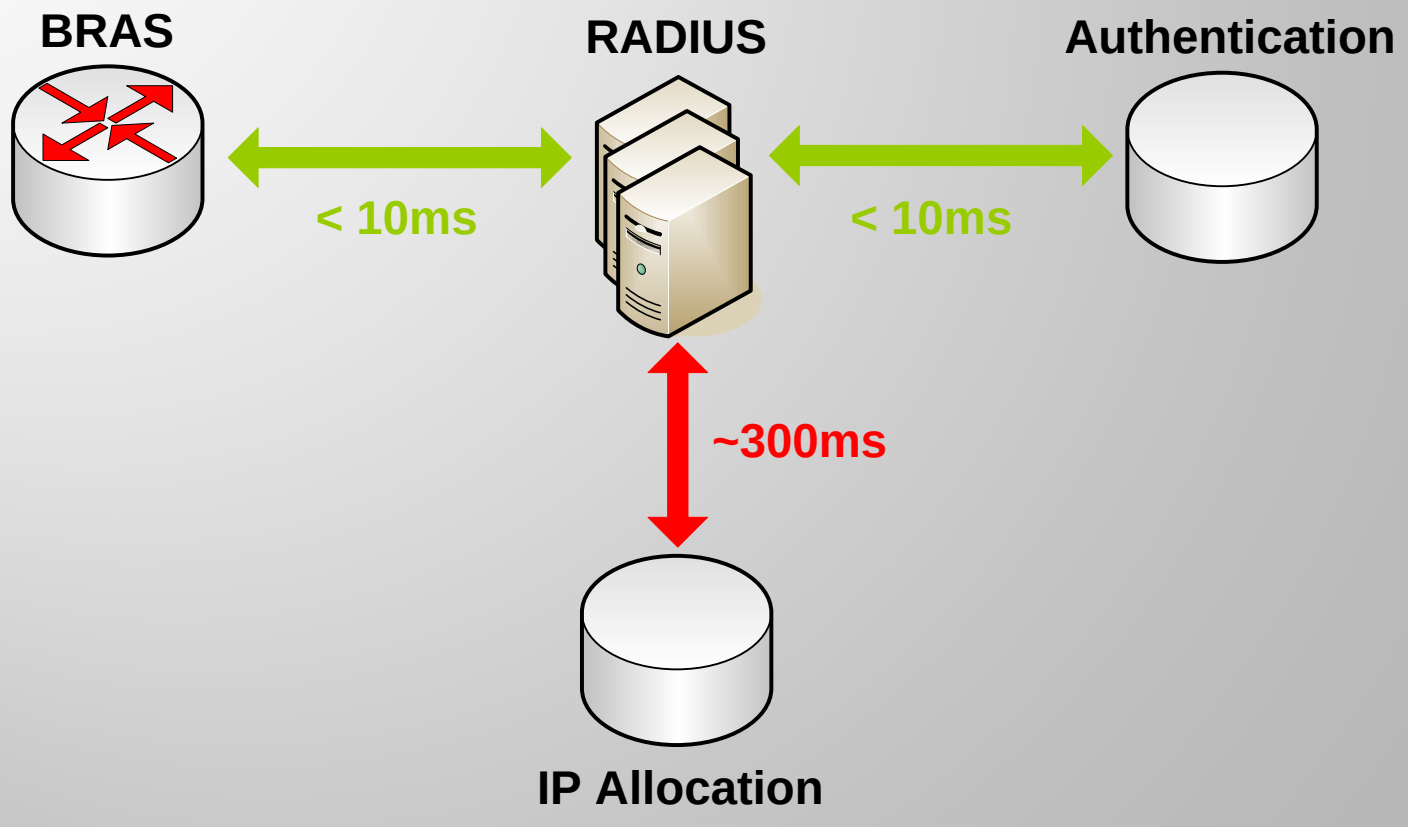

Authentication

- Solution: Local IP pools → authentication <20ms

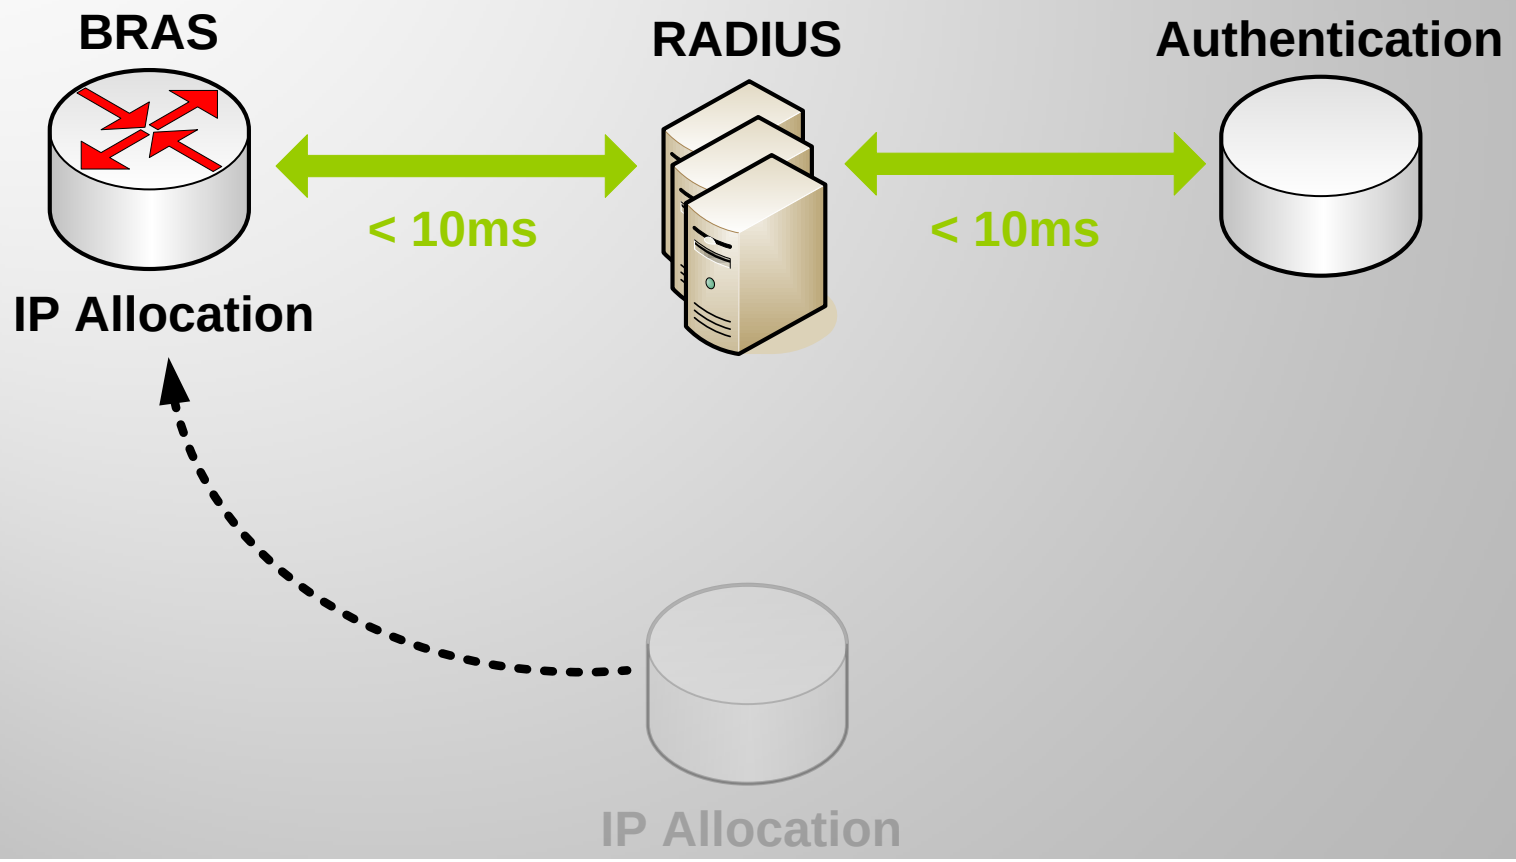

Authentication

- Goal of authenticating all users in 5 minutes achieved
  - ✓ Reduce belligerent auth requests
  - ✓ RADIUS transaction efficiency
  - ✓ Extended radius source ports
  - ✓ Local address assignment

Agenda

- Scalability in the context of.....
  - Authentication
  - **Address Allocation**
  - Accounting
  - Bandwidth

Address Allocation

- Originally, centralised allocation paradigm very efficient

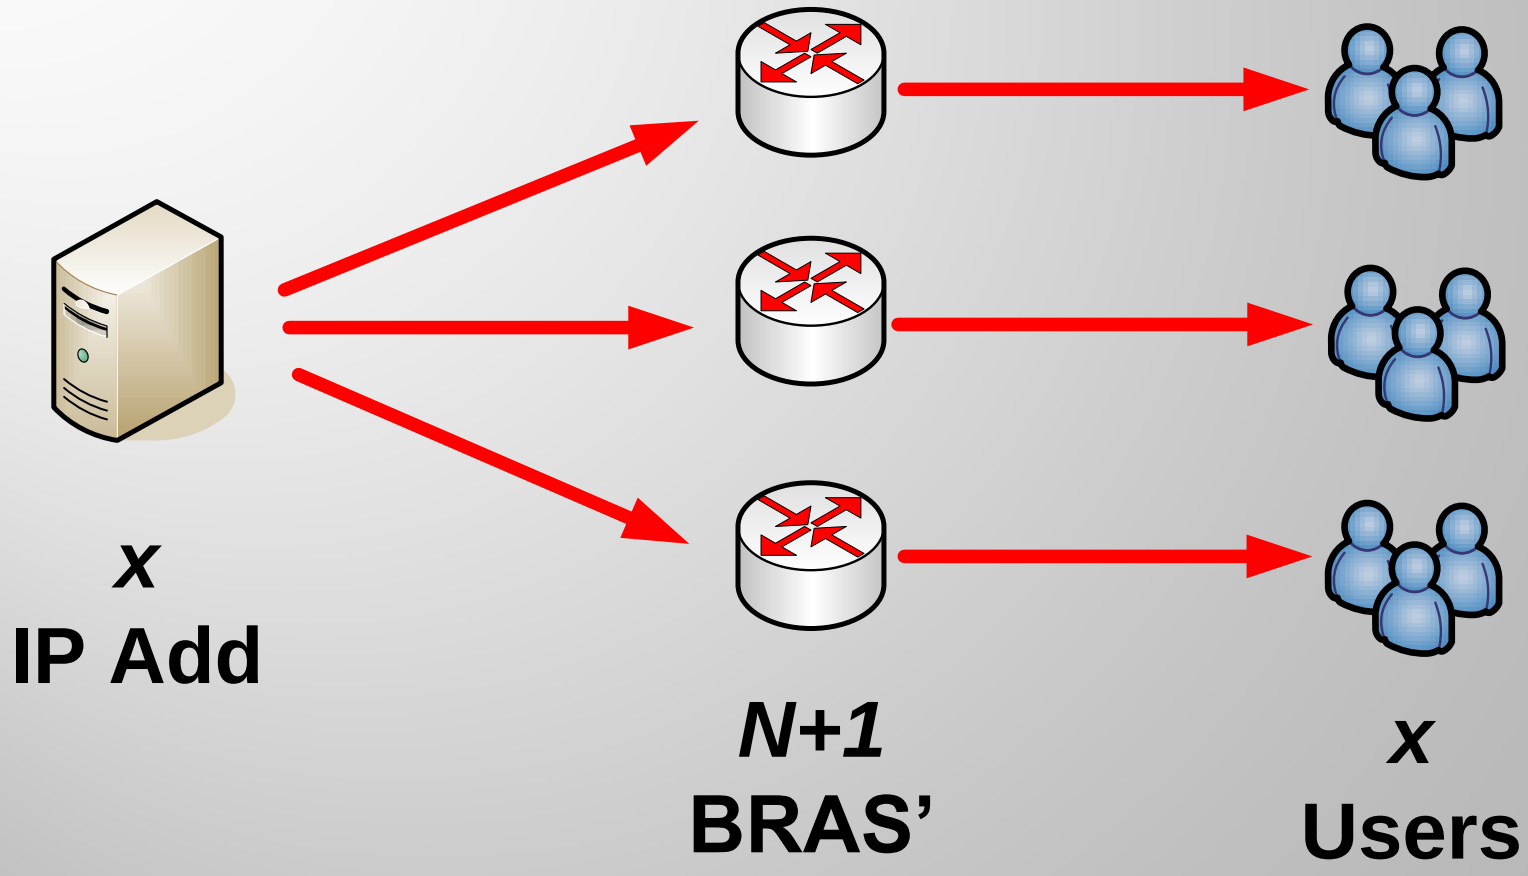

Address Allocation

- Requirements unchanged in BRAS failure scenario

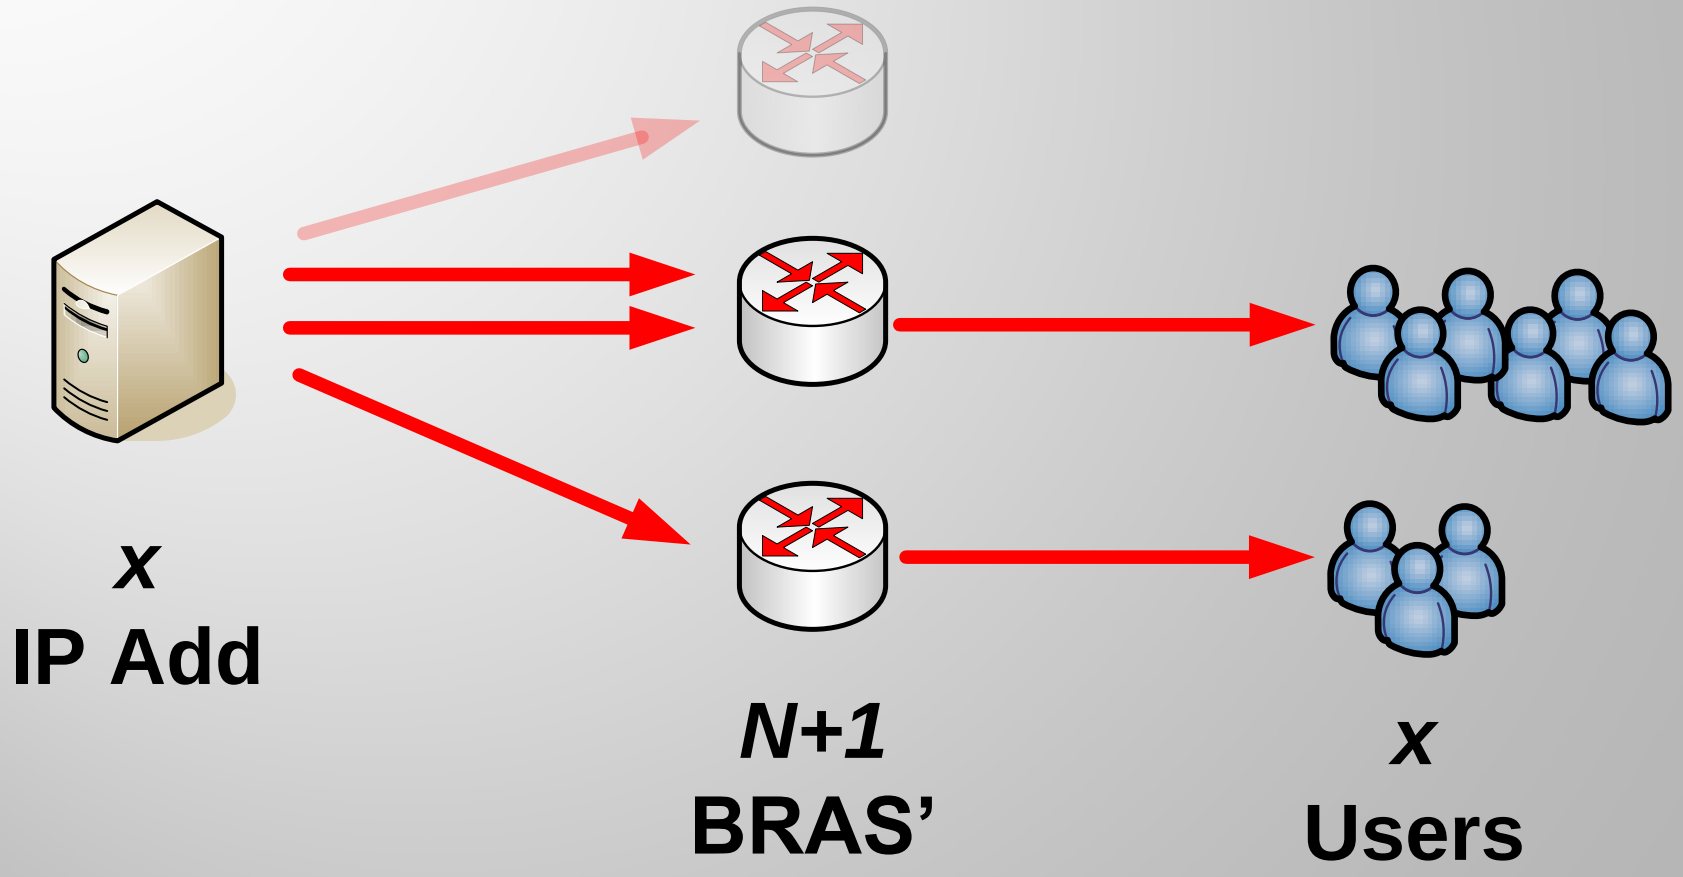

Address Allocation

- Move to local address assignment for performance

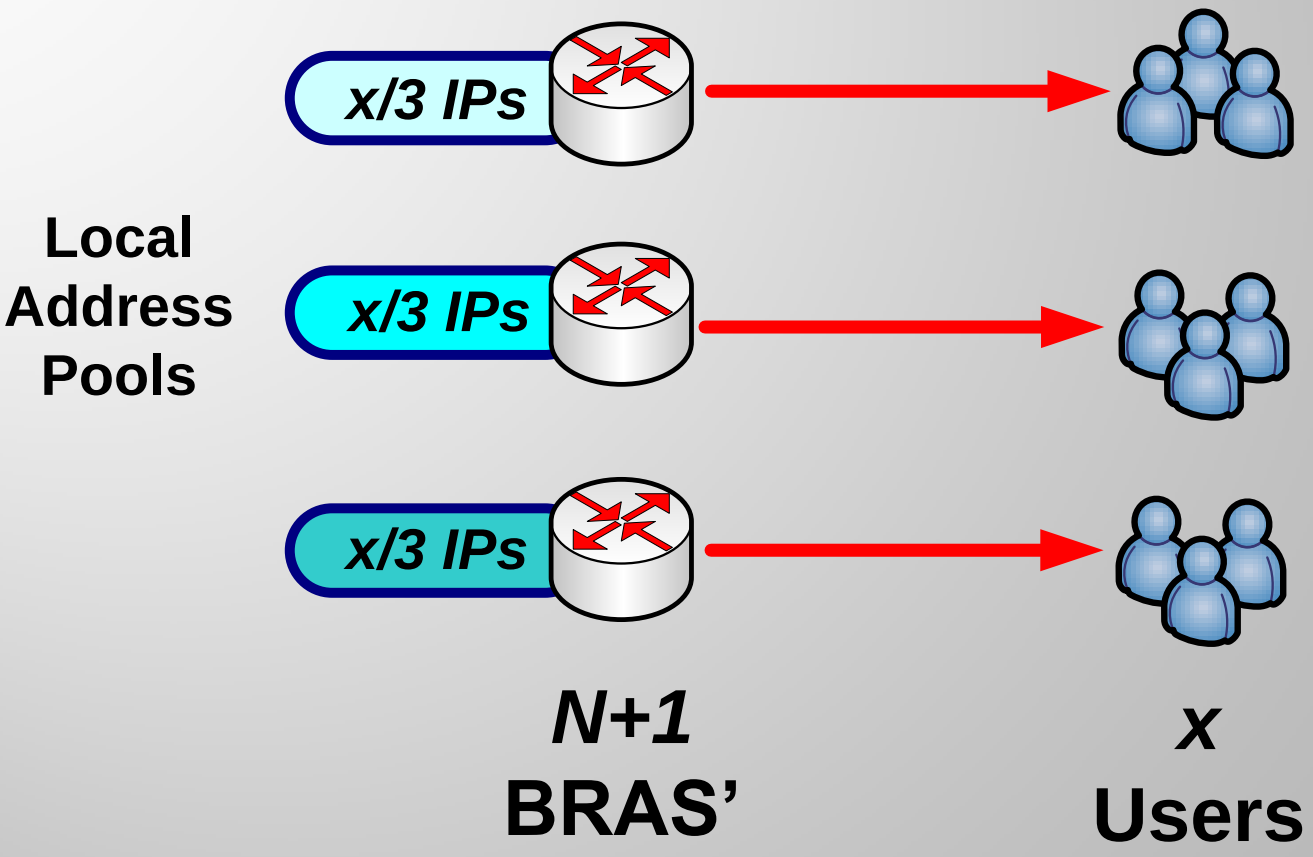

Address Allocation

- Address shortage under failure scenario

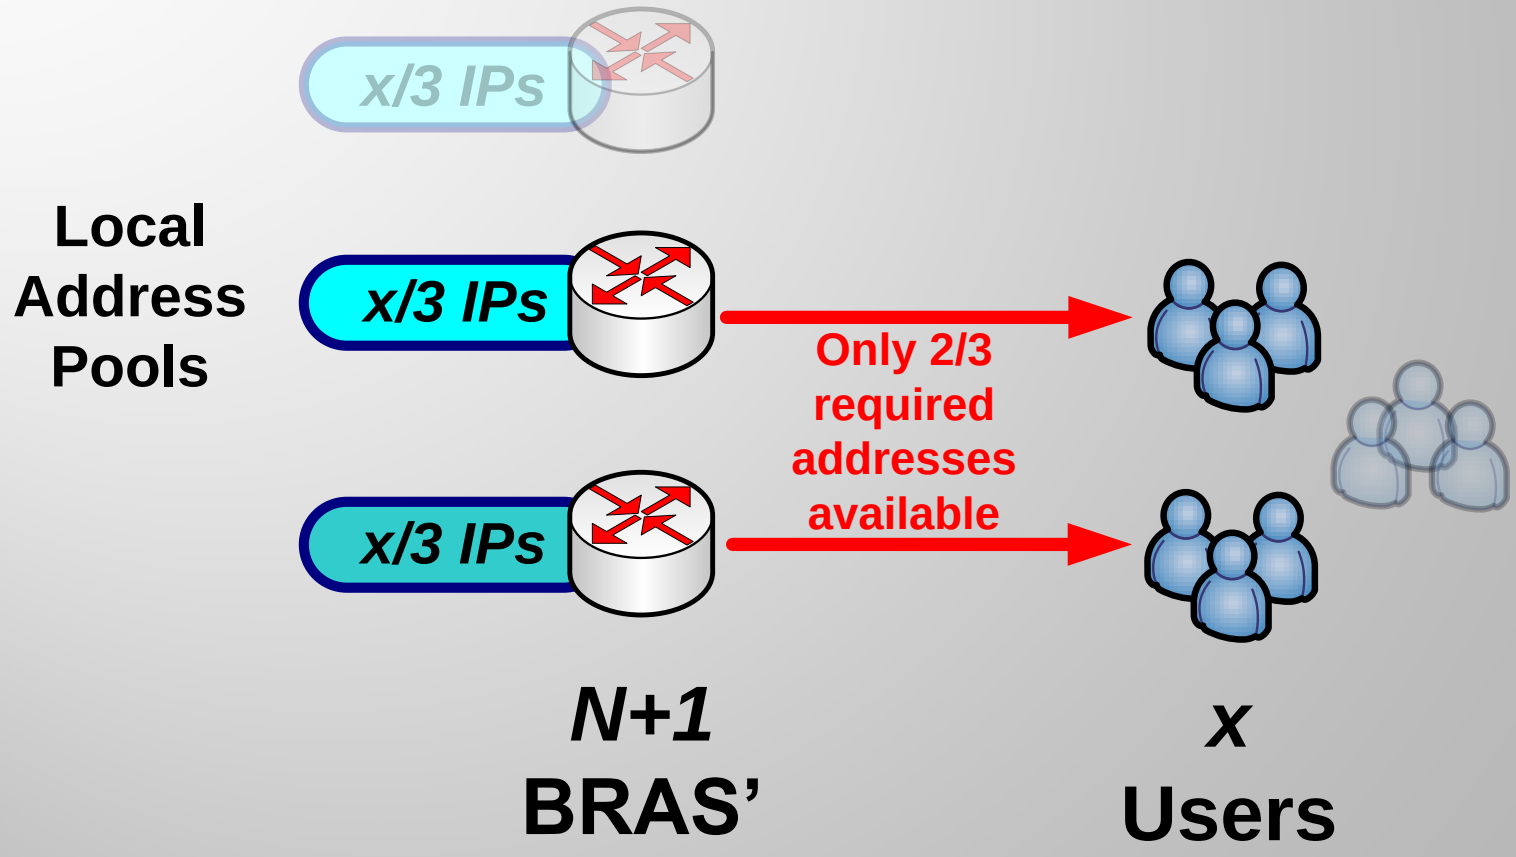

Address Allocation

- "Easy" solution – overprovision IP addresses

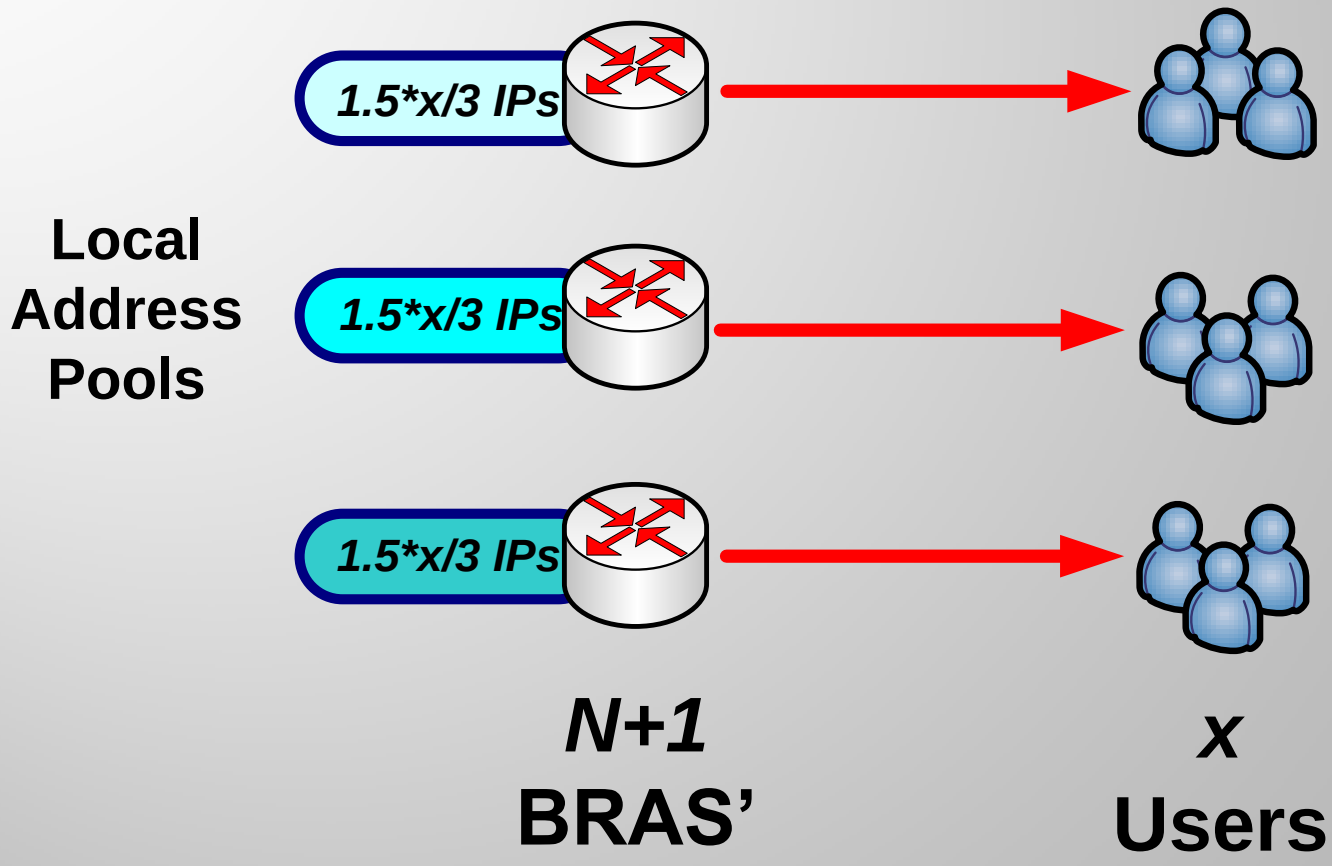

Address Allocation

- No issues in failure scenario w/ address over provisioning

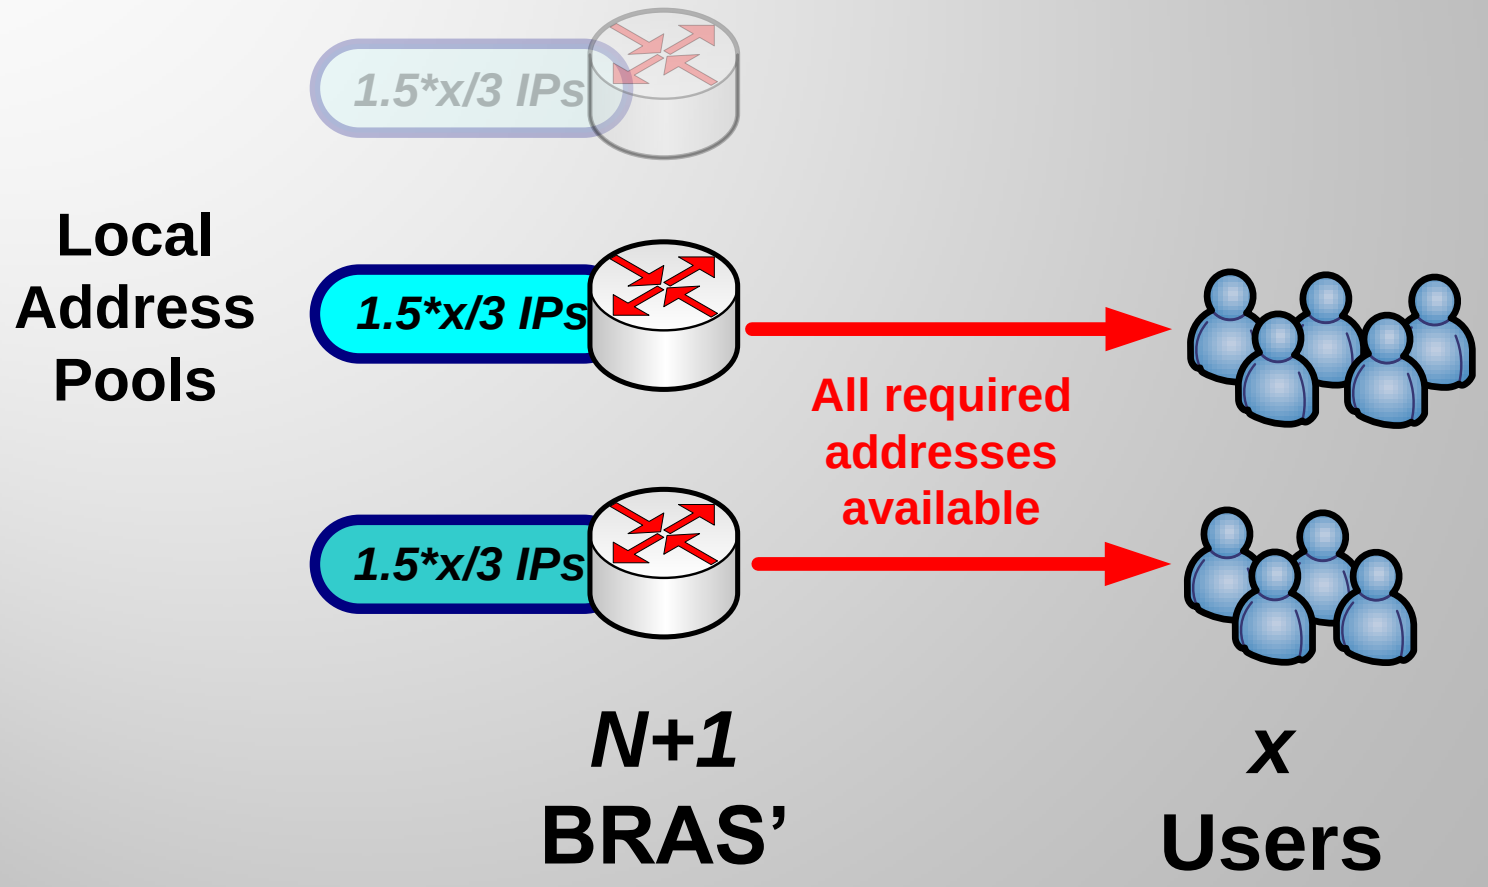

Address Allocation

- Currently trialling on-demand address pools

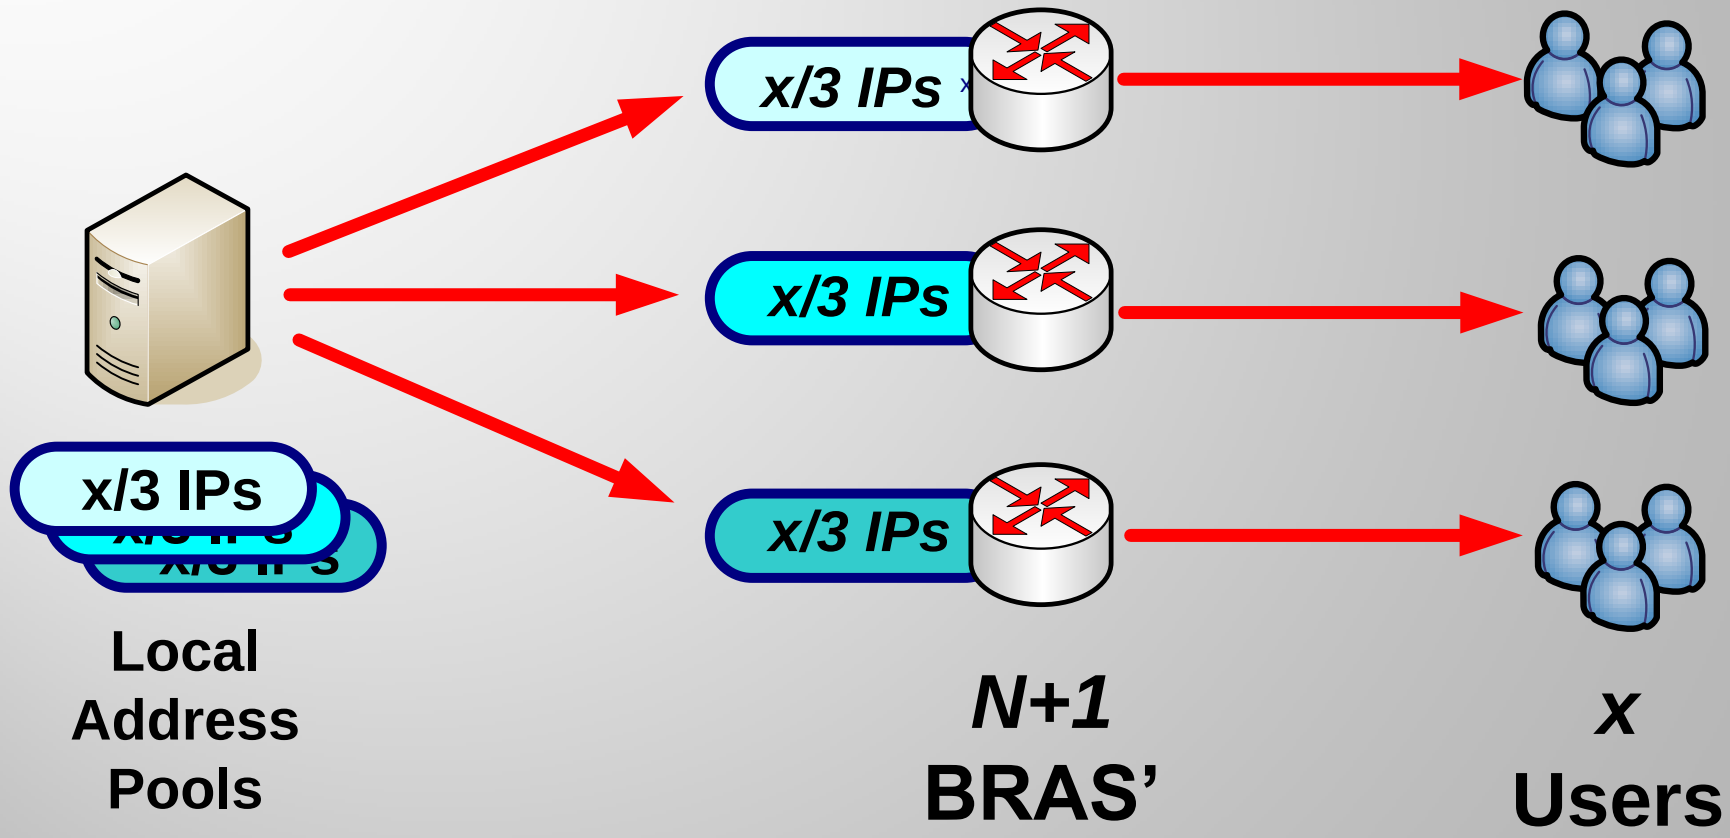

Address Allocation

- Efficiency of centralised management with performance of local assignment

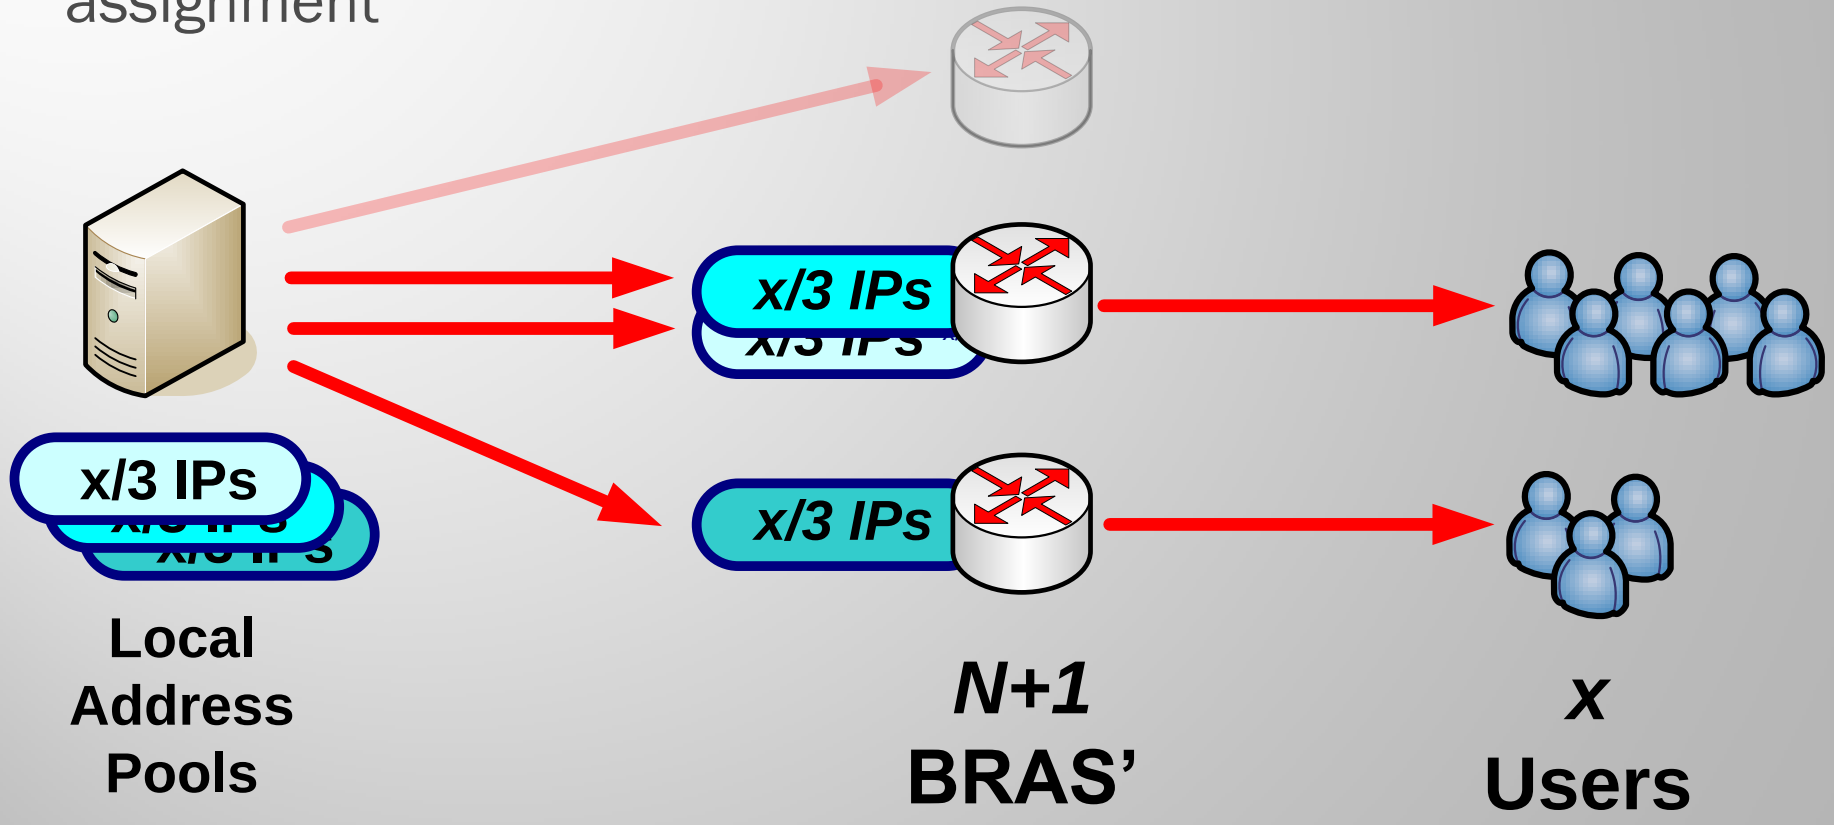

Agenda

- Scalability in the context of……
  - Authentication
  - Address Allocation
  - **Accounting**
  - Bandwidth

Accounting

- Business requirements for user traffic accounting are complex
- Traditional RADIUS, SNMP do not easily satisfy these requirements
- Important to business strategy

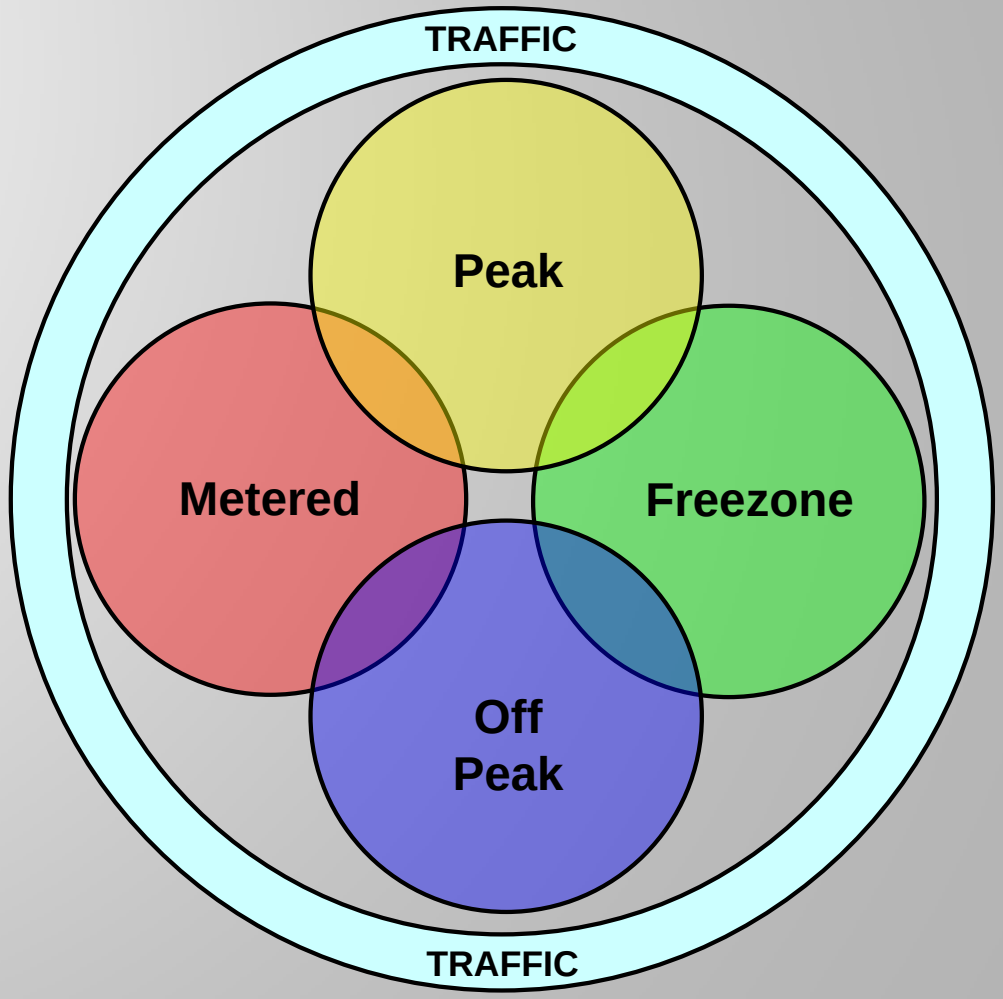

Accounting

- Traditional approach was to use Netflow accounting at the network border
- With increased DSL2+ subscribers, started to see more of this:

```
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 95%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 99%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 99%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 96%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 99%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 97%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 98%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 99%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 95%
%EARL-SP-4-NF_USAGE: Current Netflow Table Utilization is 99%
```

Accounting

- Investigated sampled Netflow…
  - Accurate for long-lived flows at high bit rates
  - Residential user traffic was mainly bursty and short-lived
  - Deemed not suitable for our accounting requirements
  - Still lives in the network used for traffic analysis.

Accounting

- Migrated from centralised to distributed Netflow model

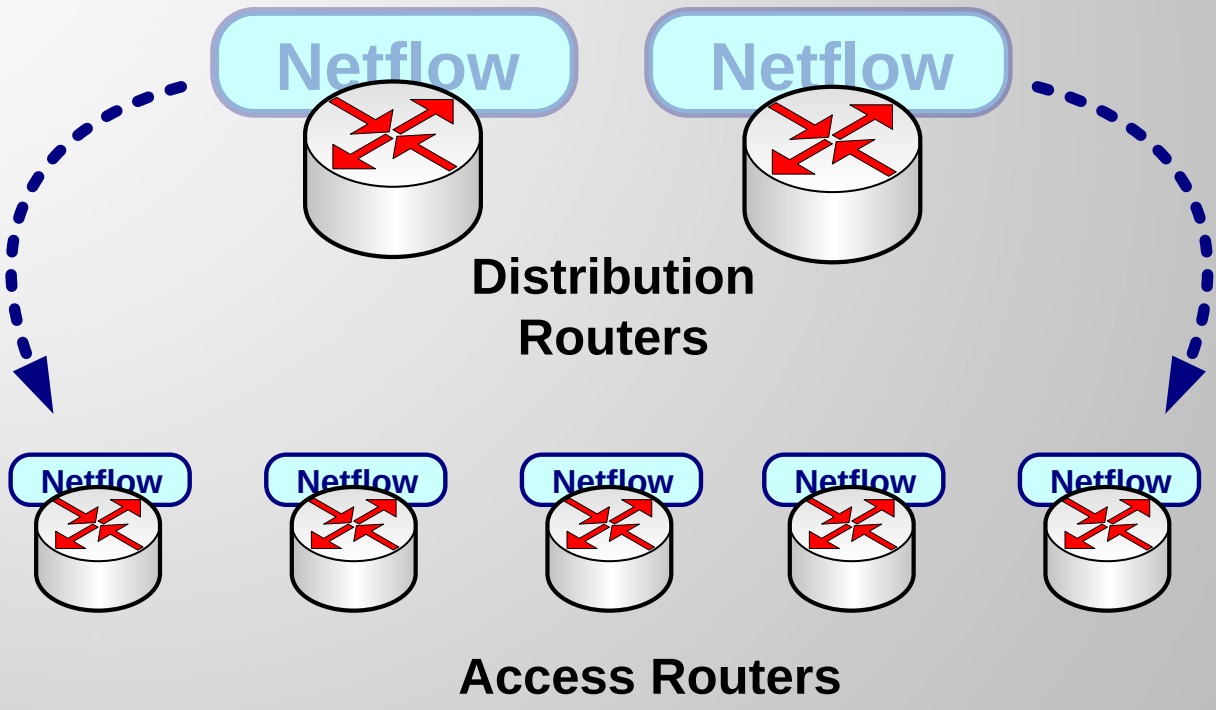

- Solved issue of losing Netflow accounting data, but introduced scaling issues with backend processing

Accounting

- Have now deployed Differentiated Radius Accounting
- Combines scalability of traditional radius with flexibility of multiple accounting buckets

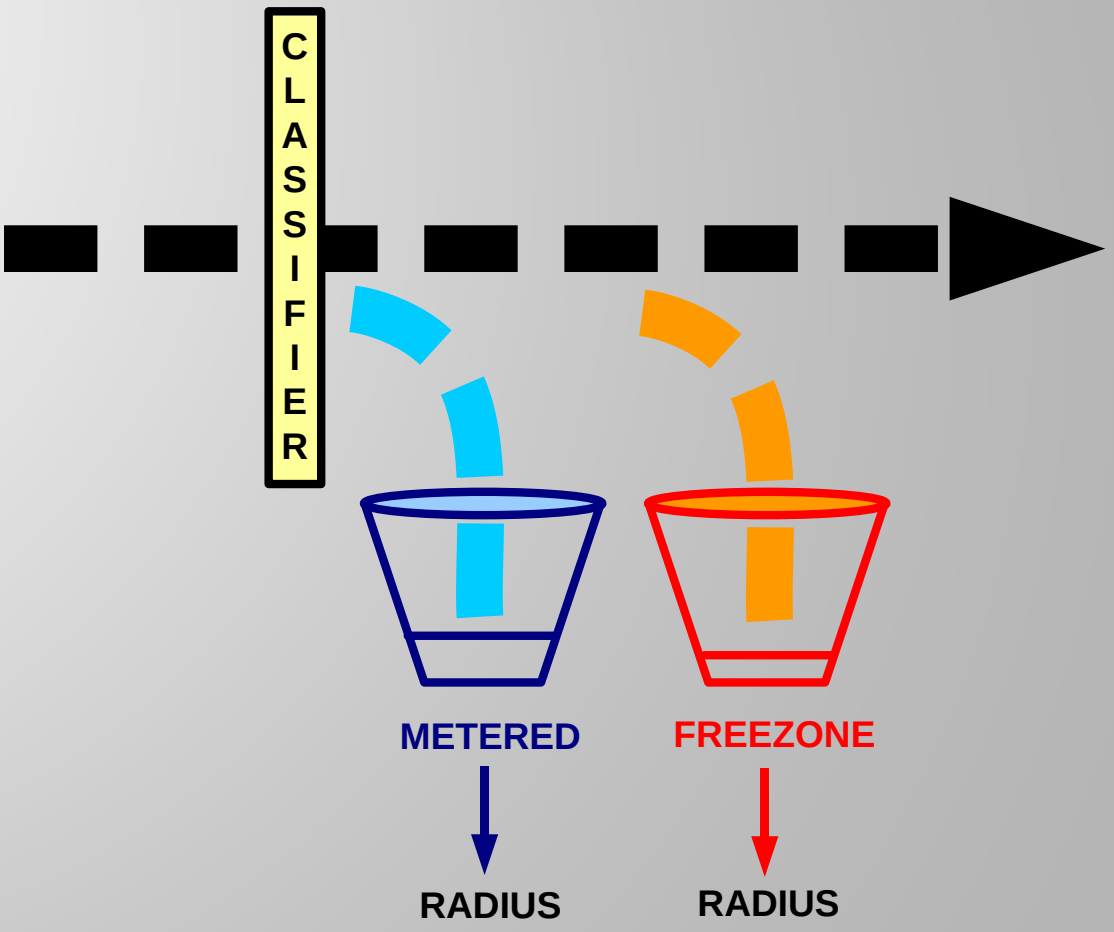

Accounting

- Differentiated RADIUS Accounting
  - Scalable – amount of accounting data does not increase as user traffic increases
  - Reduced backend system requirements
  - Accounting data visible within 15 minutes
  - Linear capacity planning - deterministic

Agenda

- Scalability in the context of……
  - Authentication
  - Address Allocation
  - Accounting
  - **Bandwidth**

Scaling Bandwidth

- Bandwidth is one of the largest cost elements of any SP network
- 2 approaches
    - Source additional capacity
    - Use what you have more efficiently
- 2 key elements to bandwidth efficiency
    - Traffic Engineering
    - Differentiated Services

Scaling Bandwidth

- Key deliverable
  - Traffic is always evenly balanced over transmission links
  - Network utilisation is unaffected by external topology changes
- Solution
  - MPLS Traffic Engineering
  - Leverage existing MPLS deployment in backbone

Scaling Bandwidth – TE

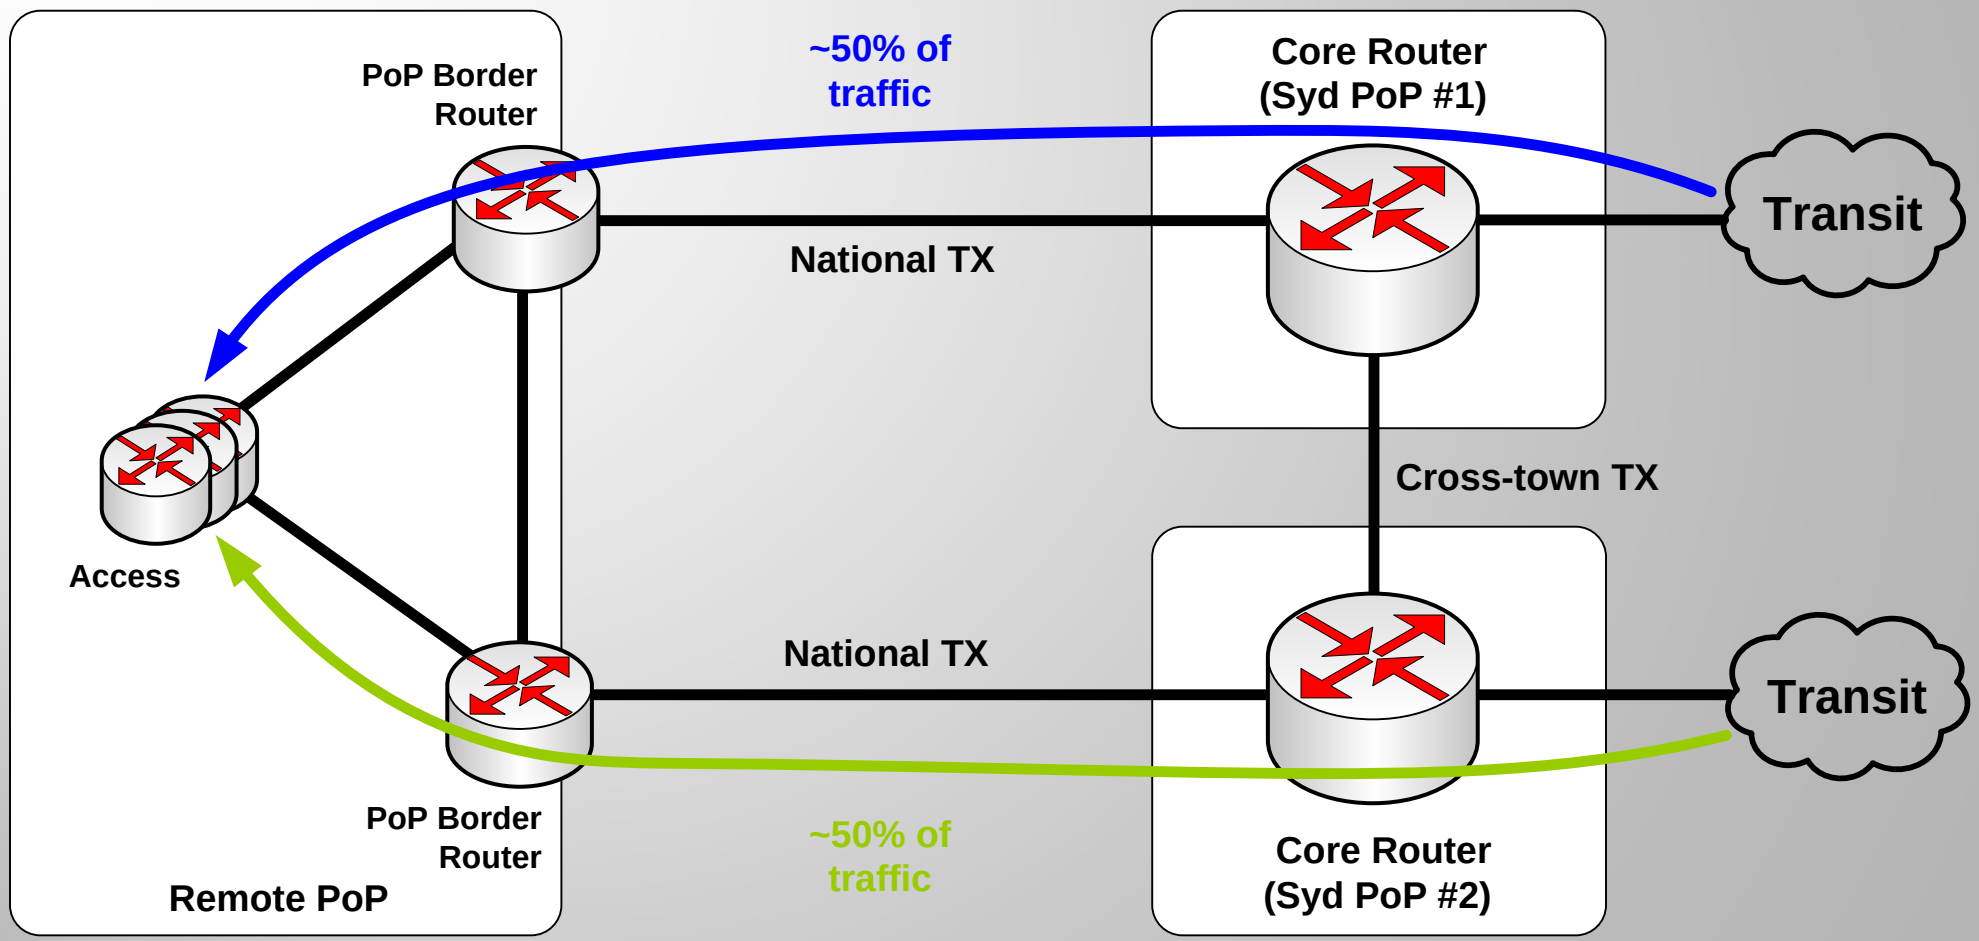

Scaling Bandwidth – TE

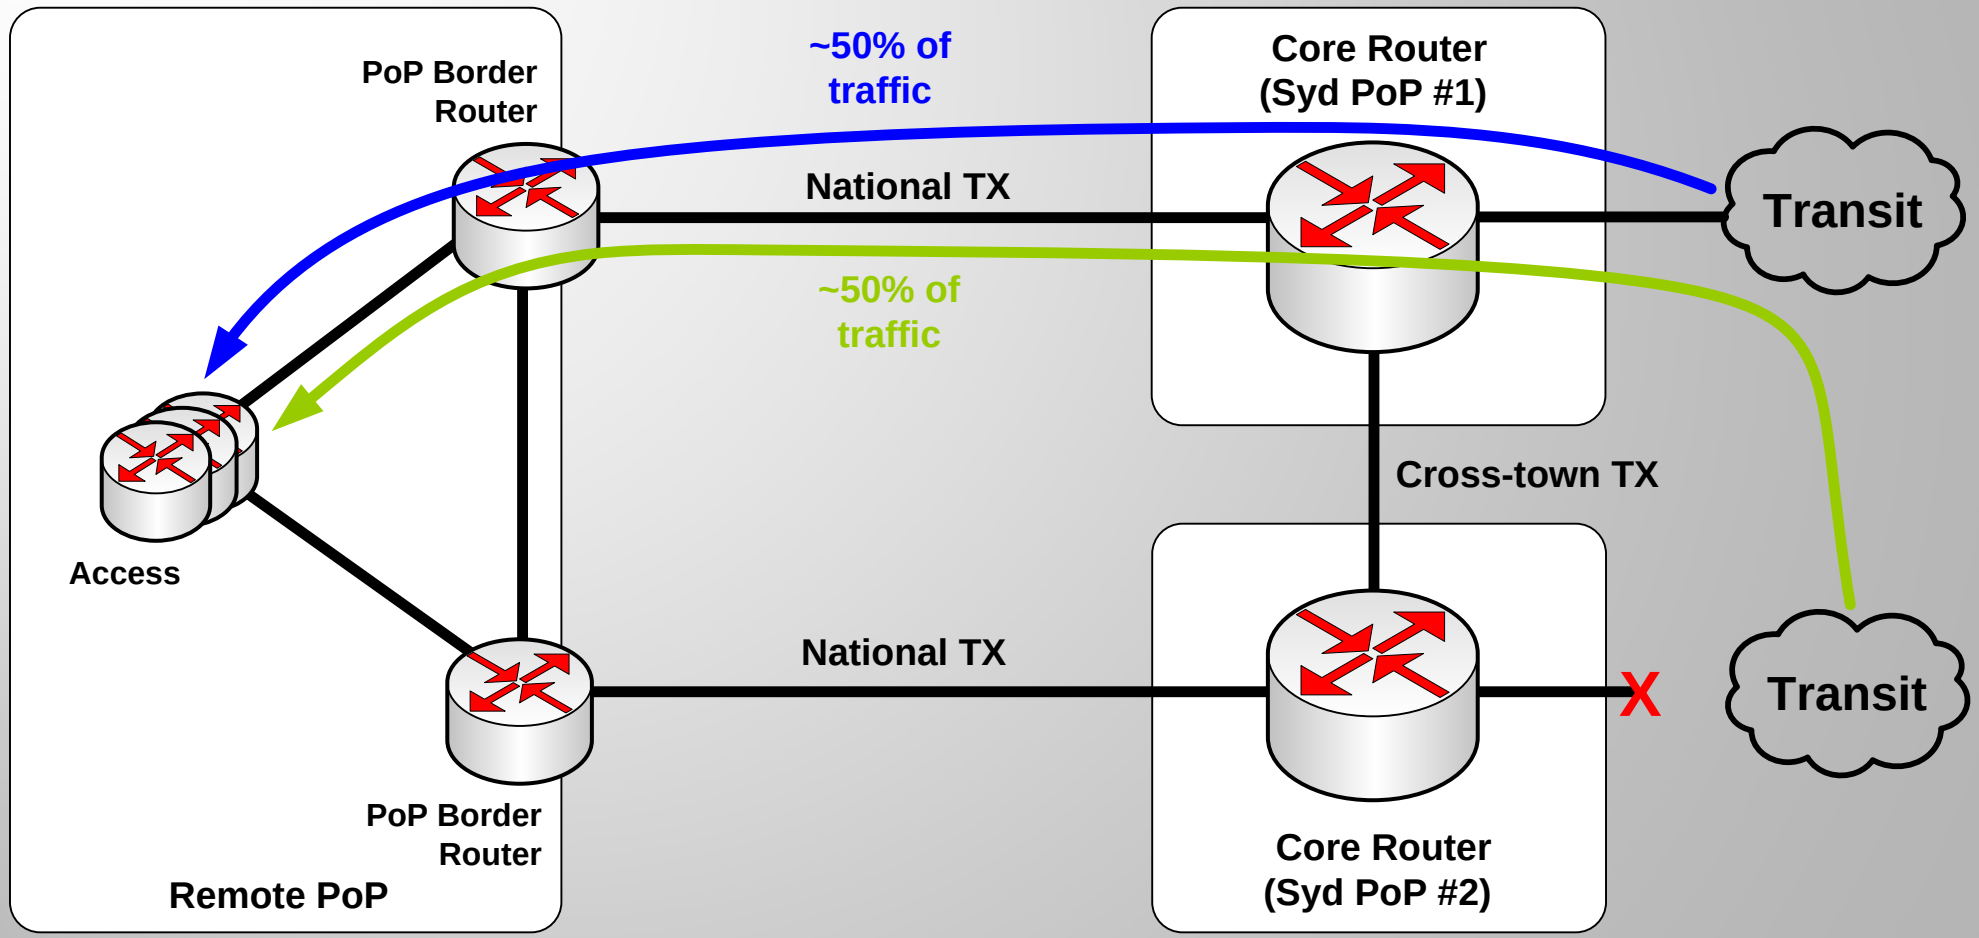

Scaling Bandwidth – TE

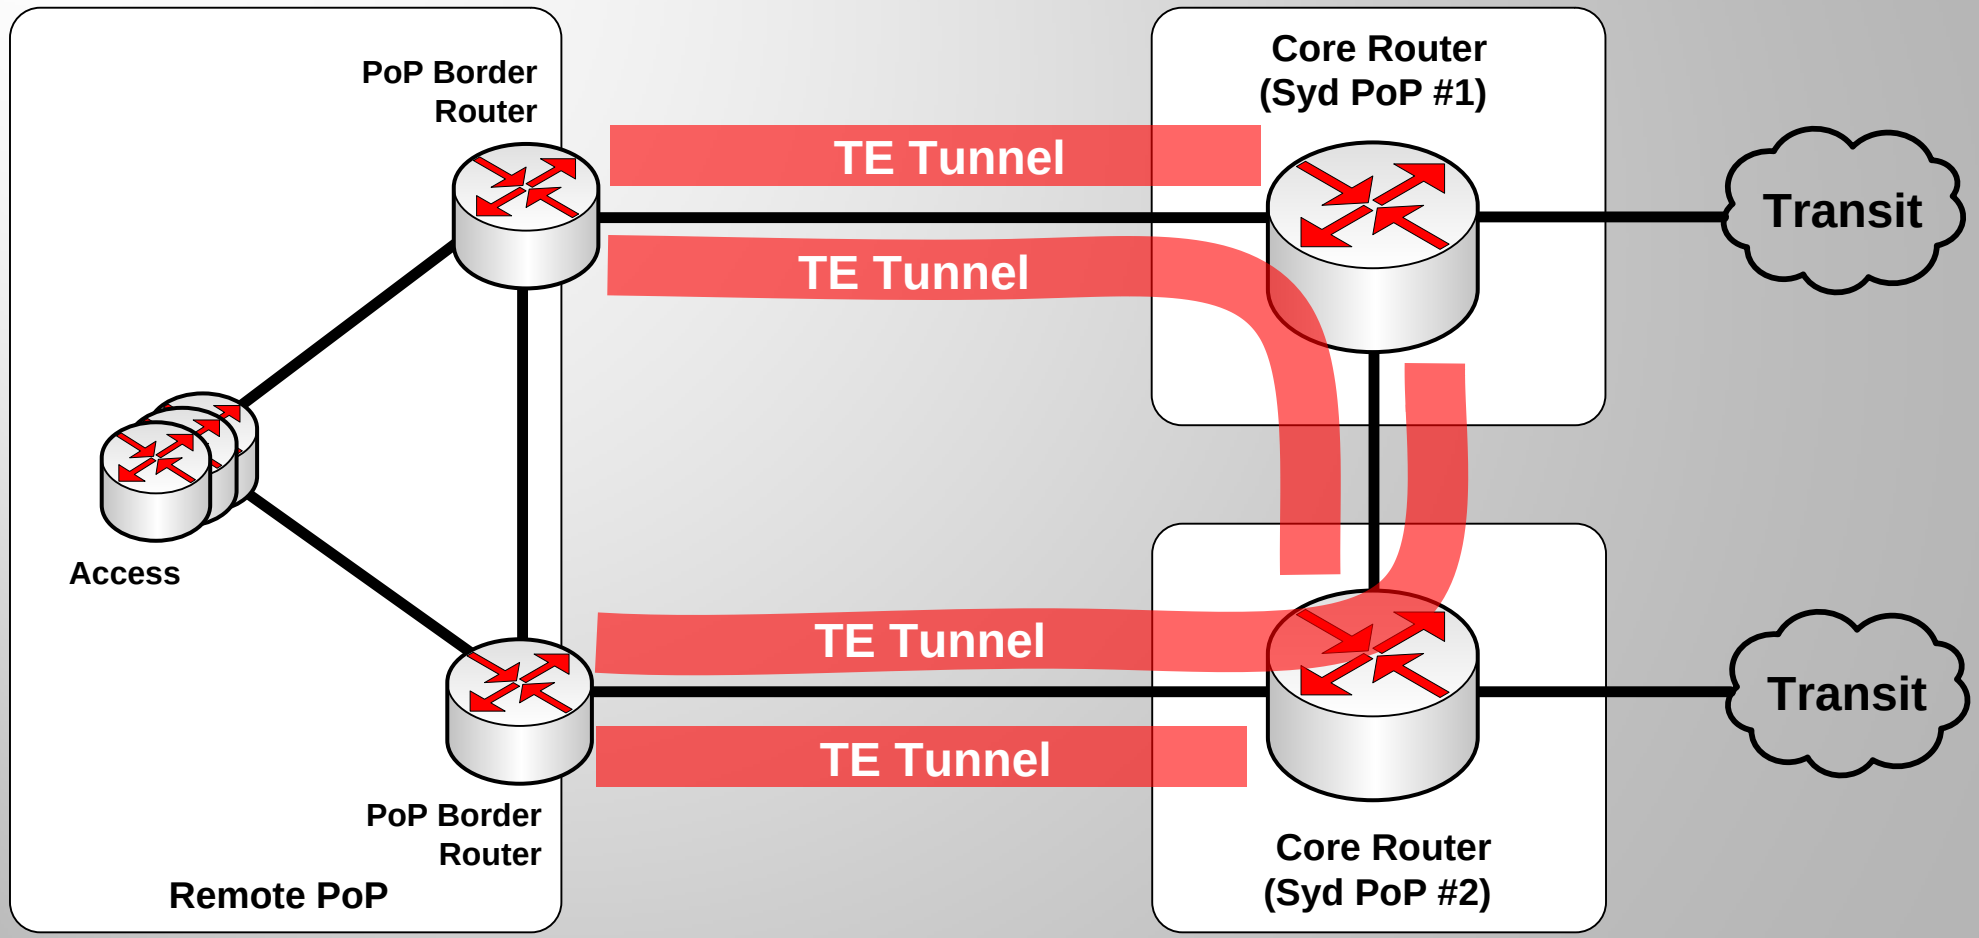

Scaling Bandwidth – TE

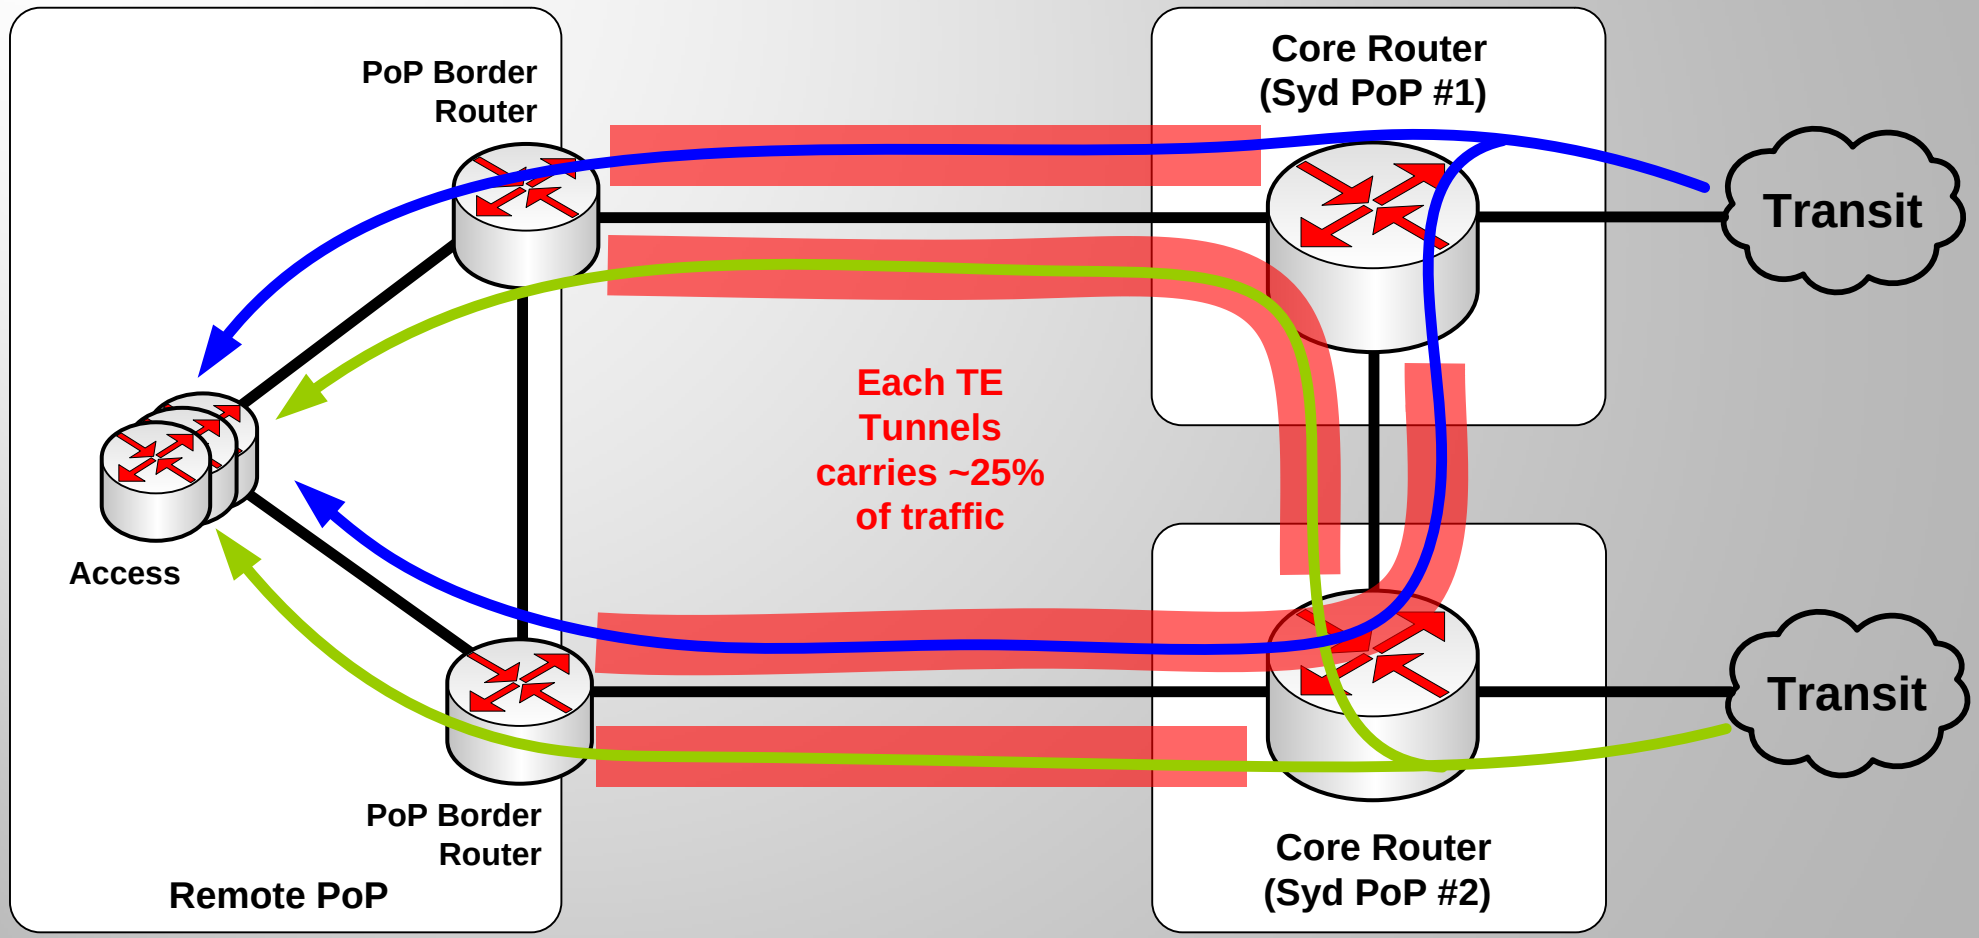

Scaling Bandwidth – TE

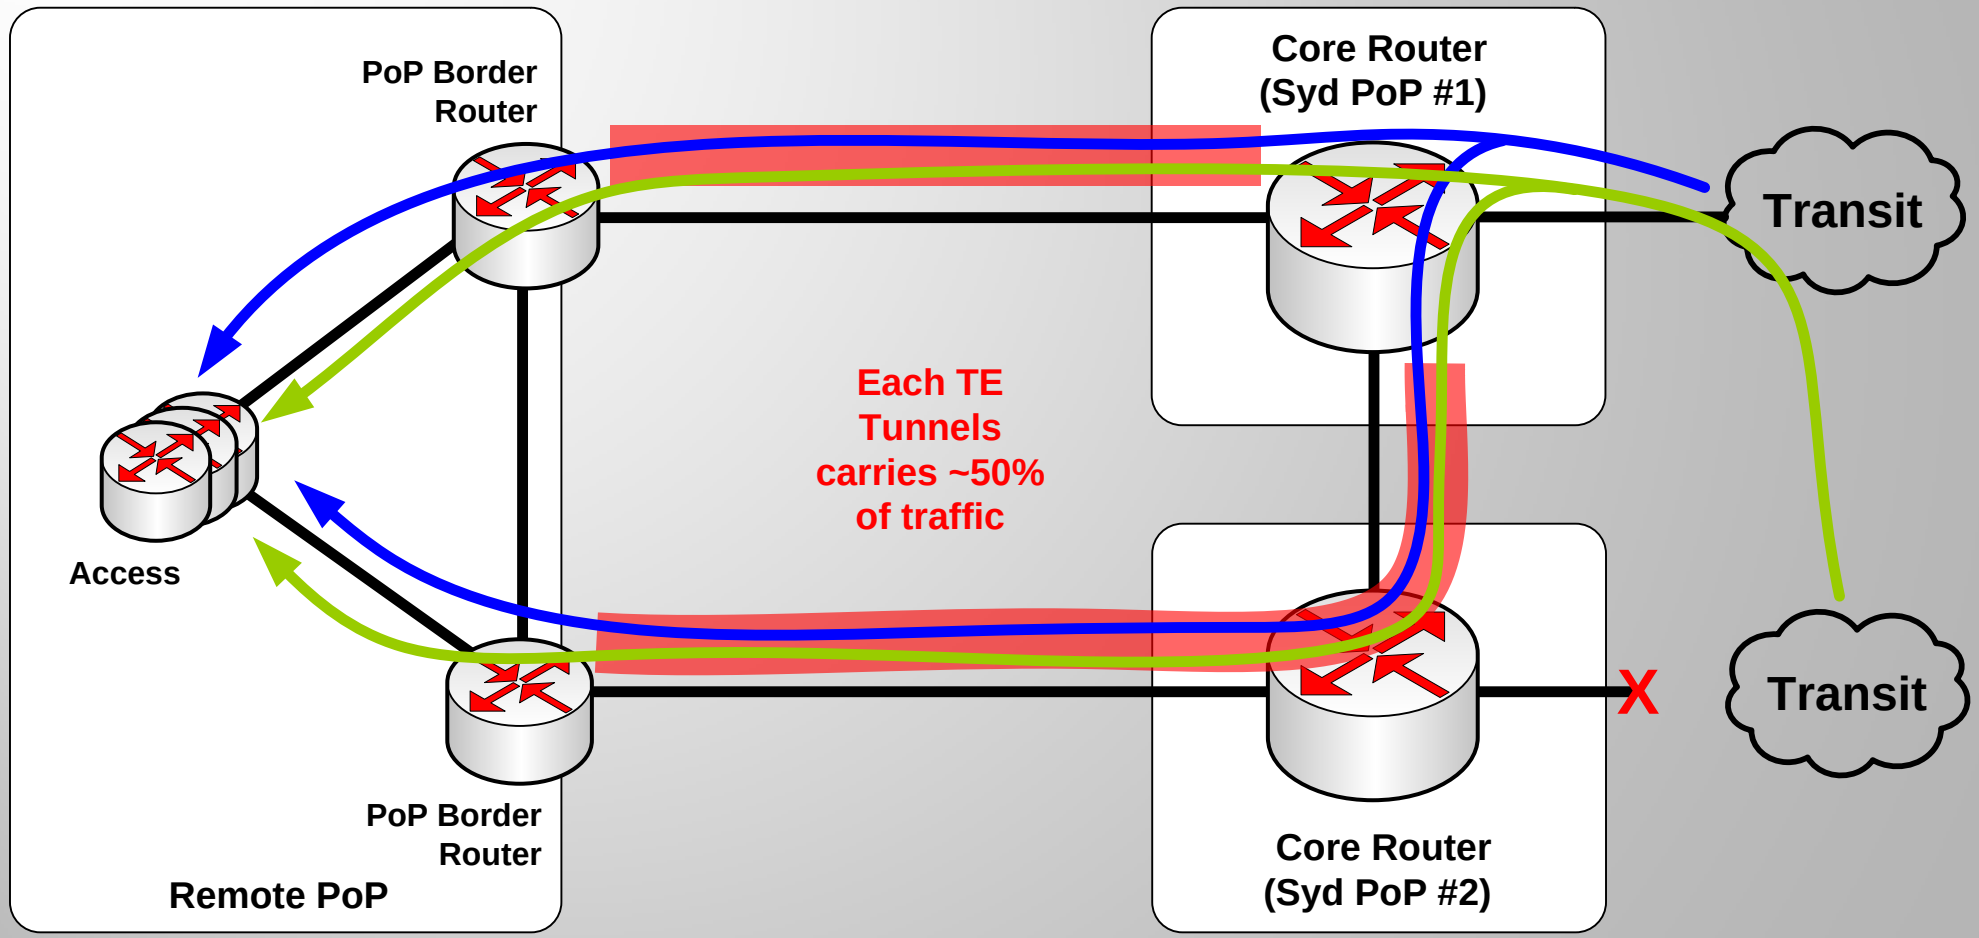

Scaling Bandwidth – TE

- Additional advantages
  - Unequal cost load sharing over multiple paths gives business flexibility
  - Fast failure recovery
- Tradeoffs
  - 50% of traffic needs to be carried by cross-town link
  - But.....plenty of capacity (30Gbps+), fibre is inexpensive

In Conclusion

- Approach to network scalability is primarily a function of deliverables vs. constraints
- Some valuable lessons learned for successful scaling
  - Network needs to be deterministic
  - Use systems that scale linearly

Thankyou