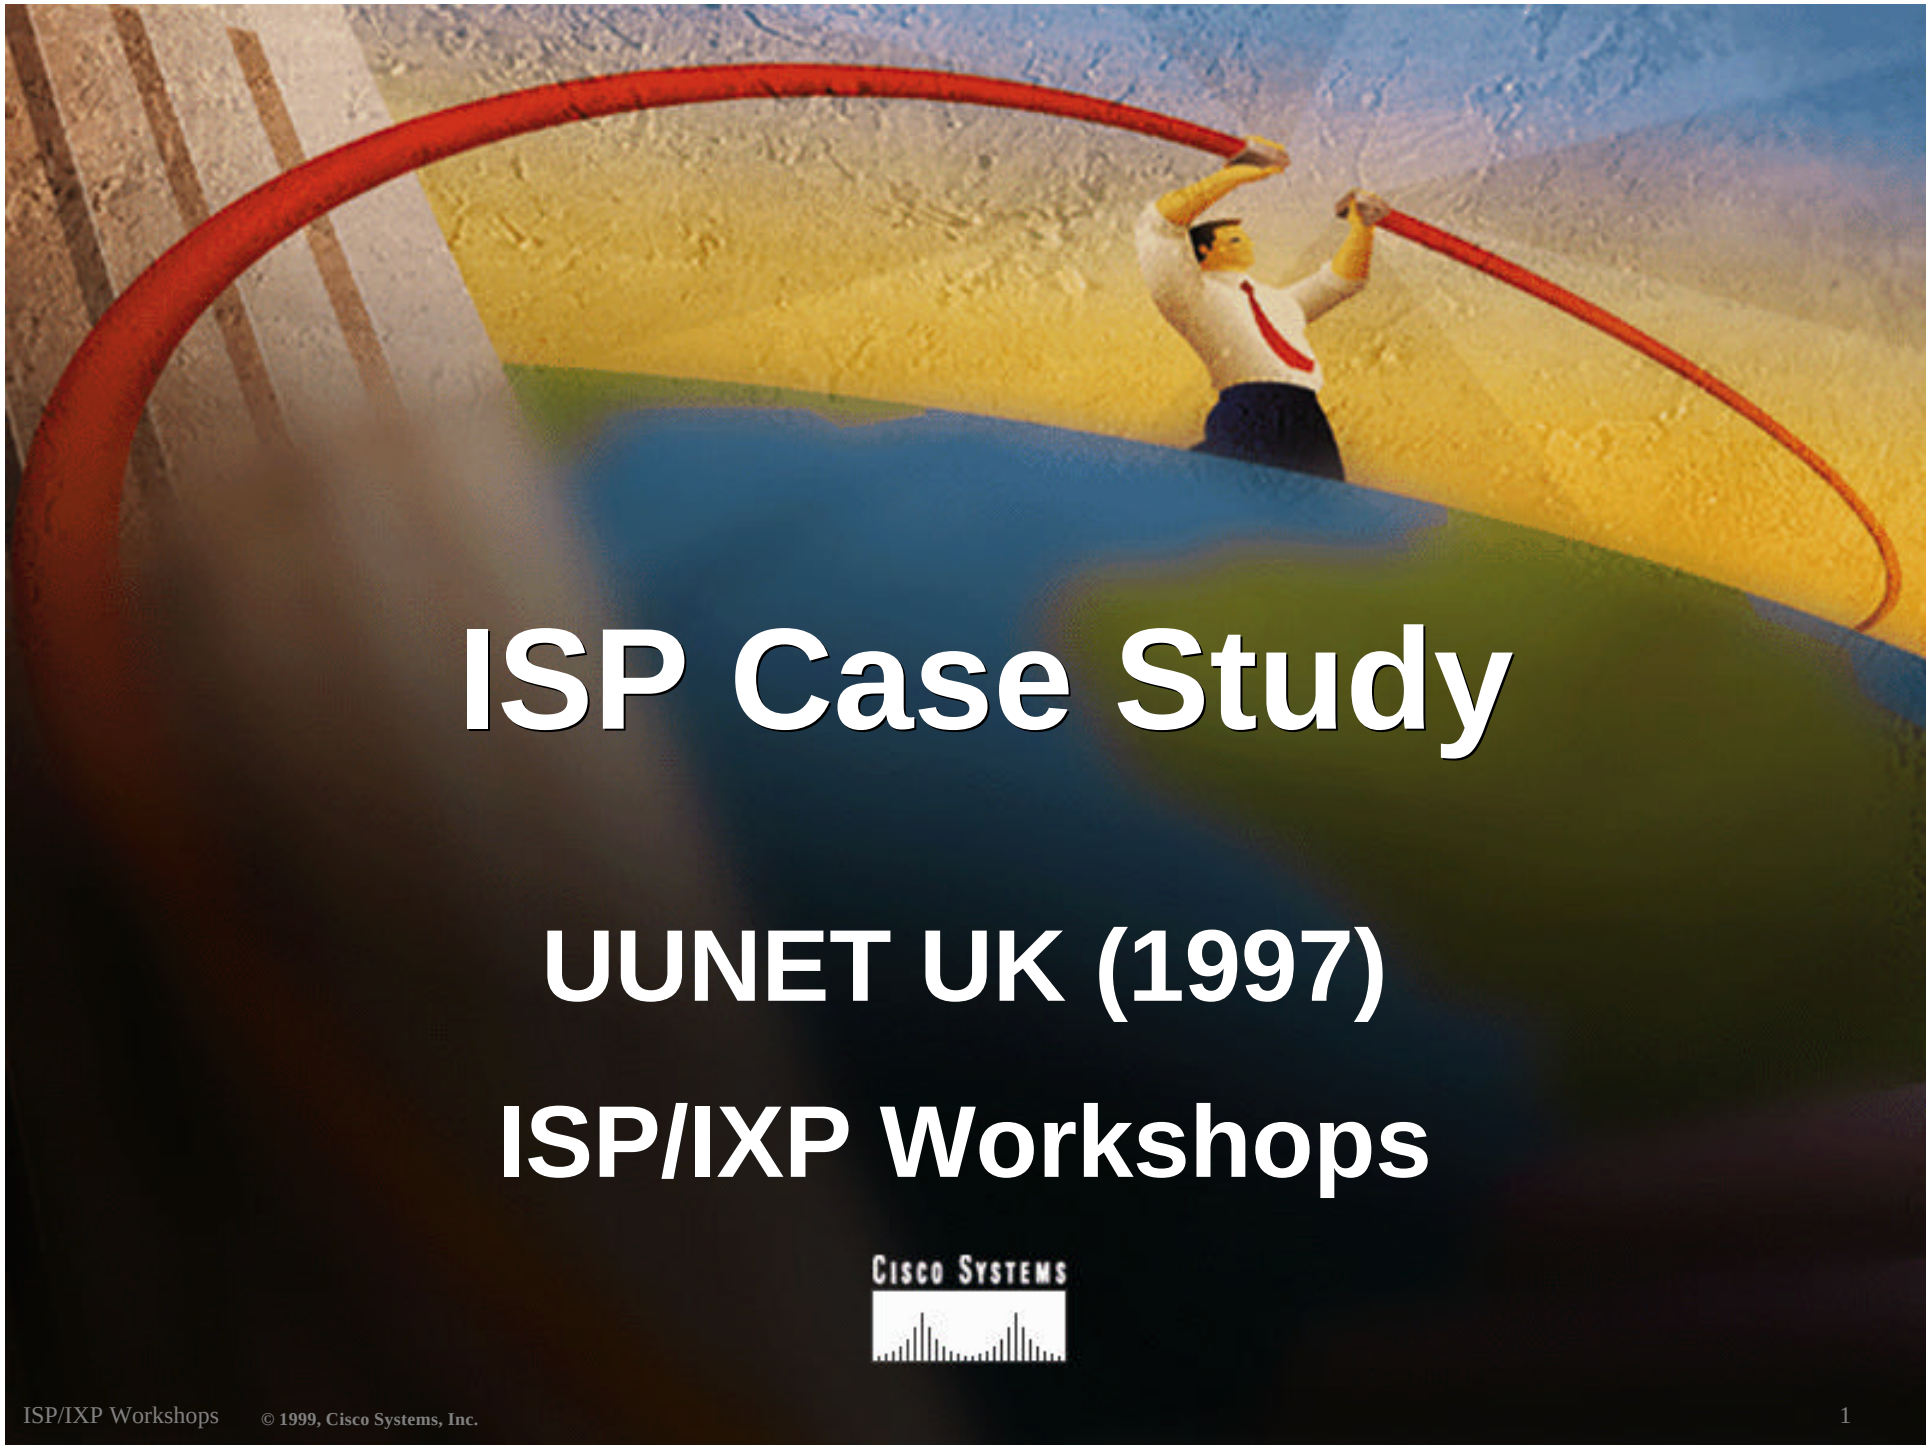

ISP Case Study

UUNET UK (1997)

ISP/IXP Workshops

Cisco Systems

Acknowledgements

- **Thanks are due to UUNET UK for allowing the use of their configuration information and network design to be made available for this presentation.**

Introduction

- History
- Network Design Principle
- PoP Design Principle
- IGP Configuration
- BGP Configuration
- Miscellaneous Features - IOS Essentials
- IP Addressing
- Systems

History

- **Founded as PIPEX in January 1992**
  - **UK's first commercial ISP**
  - **Parent company sold FTP's TCP/IP software**
- **Merged with parent company and floated on the London Stock Market**
- **European Operation**
  - **PIPEX International**
- **Bought by UUNET in 1995**
  - **now called UUNET UK**
  - **International division now UUNET Europe**

Network Design Principle

- **PoPs built in major population centres**
- **Minimum level of customer connectivity required for economic case**
- **Network borders in south of England**

  ***minimum* of two exit points**

Network Design Principle

- **Leased line backbone**

  not fully meshed

  at least triangulated

  two independent exit paths per PoP

- **Multiple line providers**

  British Telecom major provider

  also C&W, Electricity Companies, Worldcom, Scottish Telecom

National Network Layout

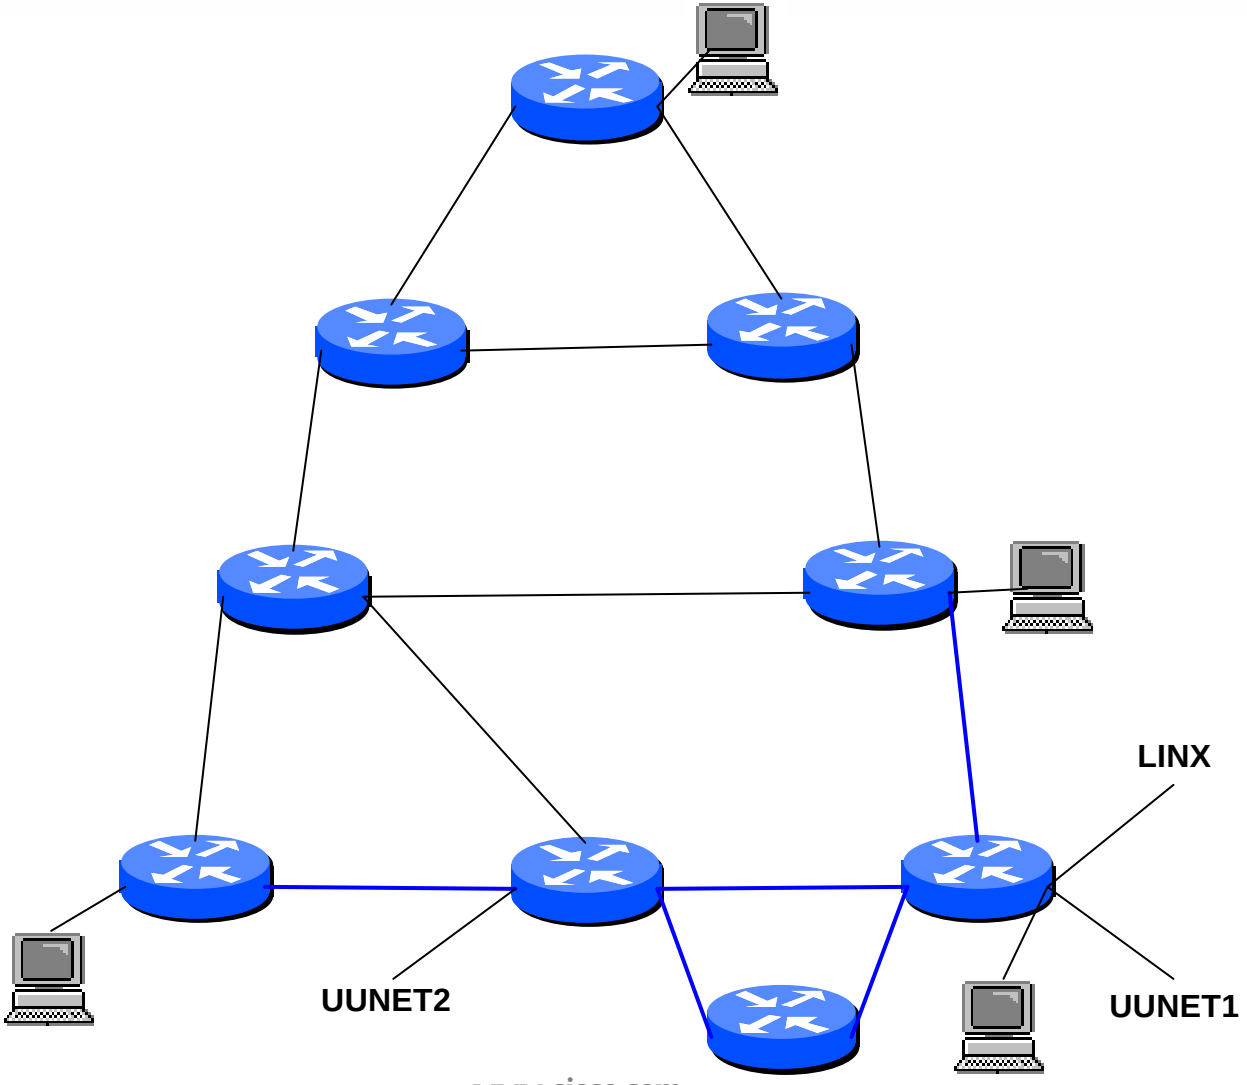

PoP Design Principles

- Core routers - carry backbone links only
- Gateway routers - carry customer links for aggregation on to backbone
- Service routers - carry LANs of hosted servers and access network
- Border routers - links to other service providers and IXPs
- Access routers - dialup customers

Typical Small PoP Design

- **Two core routers**
  - **currently 7507s with VIP interfaces**
  - **server and access LAN**
- **One gateway router**
  - **currently 7507 with two FSIP8, two MIP2**
- **Console server**
  - **out of band access for PoP equipment**
  - **modem dialup**
- **FDDI ring**

Small PoP Layout

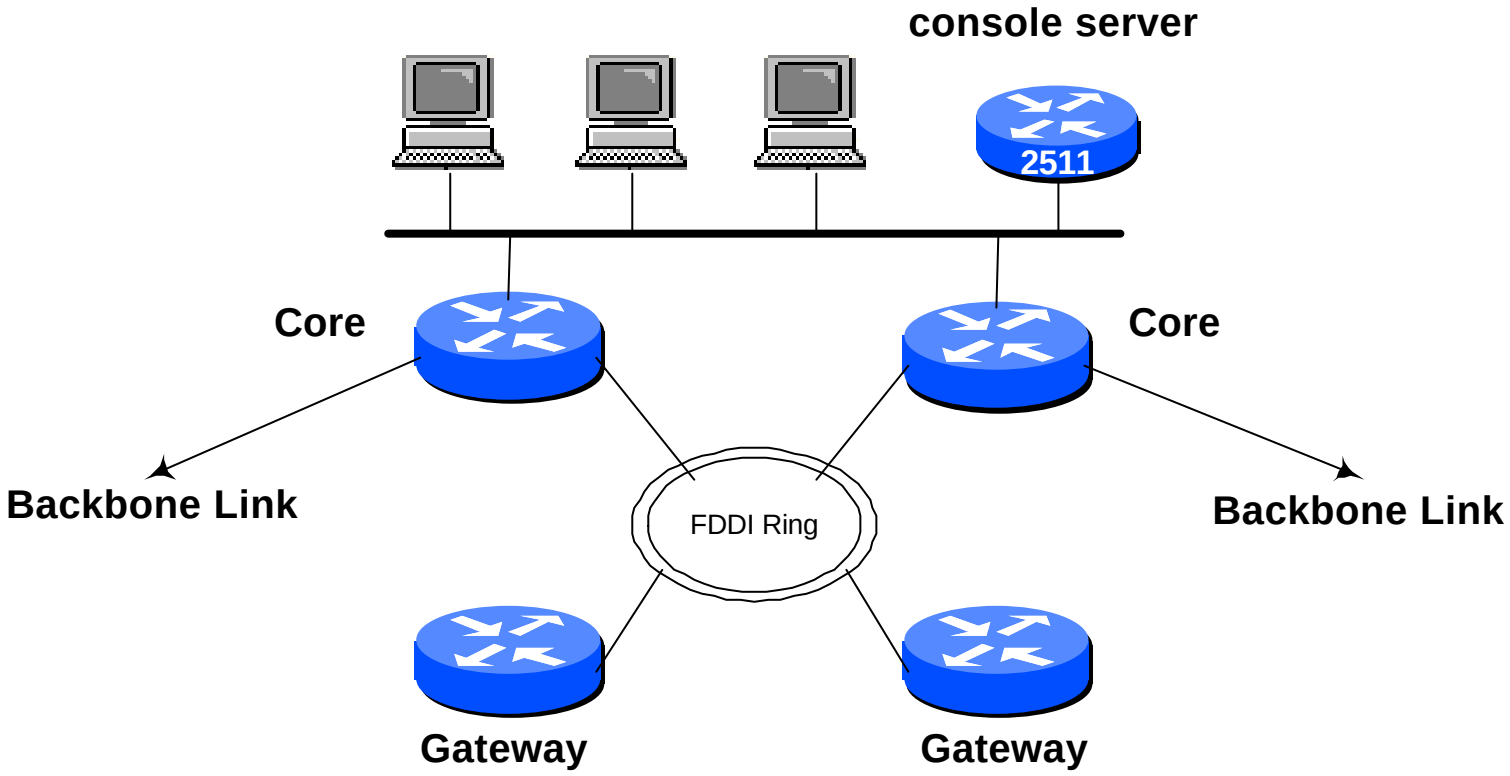

Typical Large PoP Design

- **Two core routers**
  - **currently 7507s with VIP interfaces**
  - **server LAN**
- **Two or more gateway routers**
  - **currently 7507 with two FSIP8, two MIP2**
- **Two services routers**
  - **currently 7507 with VIP or *x*IP interfaces**
  - **hosted services, access network**

Typical Large PoP Design

- **Border Router**
  - **currently 7507 with RSP4 and 256Mbytes**
- **Console servers**
  - **out of band access**
  - **modem dialup**
- **Catalyst 5000 switched backbone**
  - **dual switches**
  - **VLANs**
  - **fastetherchannel**

Large PoP Layout

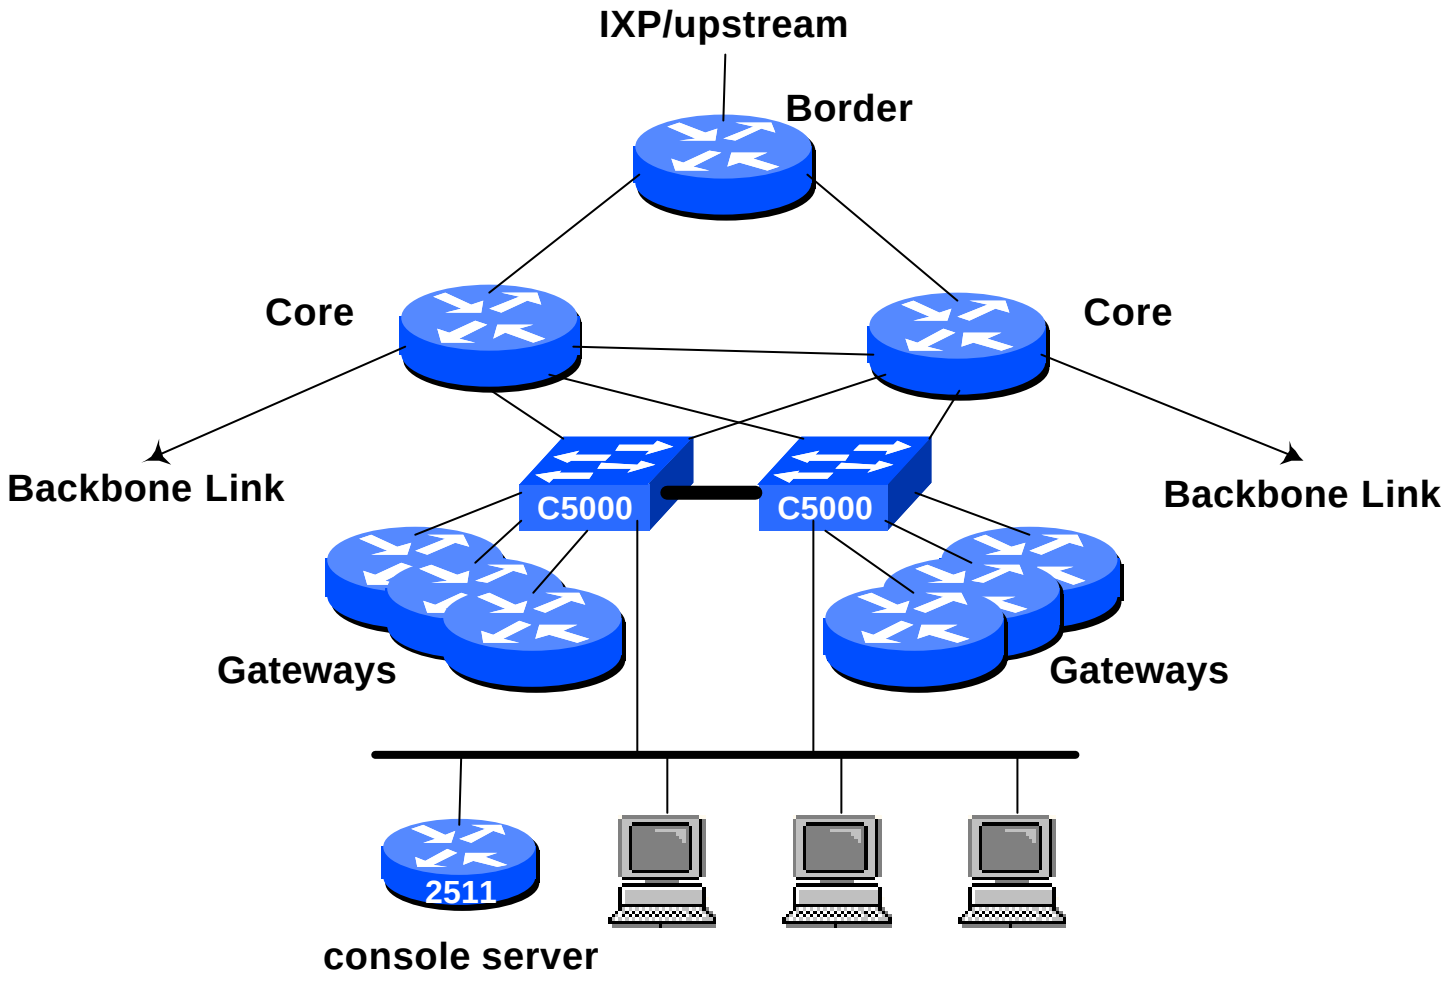

IGP Configuration

- Started with IGRP, upgraded to EIGRP, migrated to OSPF
- Core backbone is OSPF area 0
- Each PoP is an OSPF area
- Networks summarised between areas
  - keeps IGP small
  - rapid convergence in case of link failure
- Design first used for EIGRP
  - nets summarised on PoP boundaries

BGP Configuration

- Each PoP is a route reflector cluster
- Core routers are route reflectors
- Remaining routers are cluster clients
- Keeps iBGP mesh small

  core routers only

  easier to add more routers to network

BGP Configuration

- **Only core routers carry full routes**
  - **other routers carry UK routes**
- **Upstream border routers**
  - **carries full routes**
- **Local/Regional border routers**
  - **carries learned external routes**
  - **carries domestic UUNET UK routes**
  - ***no*** **default route**

BGP Configuration (continued)

- Extensive use of communities to distinguish between types of customers and routes.
- Route flap dampening enabled on the edges
- Internet Routing Registry (IRR) used
  - RIPE Routing Registry
  - registering external routing policy
  - peers only accept what is in IRR
- AS and route filtering on edges!

Route Flap Dampening

basic bgp configuration and implementation of route-map

```
router bgp 1849
bgp dampening route-map expo-flap-dampen
```

no flap dampening for key user defined networks defined in access-list 189

```
route-map expo-flap-dampen deny 5
 match ip address 189
```

no flap dampening for root nameserver /24 networks in access-list 180

```
route-map expo-flap-dampen deny 7
 match ip address 180
```

flap dampening for 192/8 network block (access-list 188)

```
route-map expo-flap-dampen permit 9
 match ip address 188
 set dampening 30 750 3000 60
```

Route Flap Dampening

flap dampening for all the other /24 networks not in 192/8 netblock

```
route-map expo-flap-dampen permit 10
 match ip address 181
 set dampening 30 750 3000 60
```

flap dampening for all /22 and longer prefixes

```
route-map expo-flap-dampen permit 20
 match ip address 182
 set dampening 15 750 3000 45
```

flap dampening for all remaining prefixes

```
route-map expo-flap-dampen permit 40
 set dampening 10 1500 3000 30
```

Note that the cisco defaults are **`set dampening 15 750 2000 60`** and are what would be applied using the default dampening configuration.

BGP Communities

```
1849:70          set local pref to 70 for multihomed customers (see RFC1998)
1849:80          set local pref to 80 for multihomed customers
1849:90          set local pref to 90 for multihomed customers
1849:110         set local pref to 110 for multihomed customers
1849:130         set local pref to 130 for multihomed customers
1849:701         routes learned from UUNET USA
1849:702         routes learned from UUNET Europe
1849:703         routes learned from UUNET Asia-Pacific
1849:5000        Customers and backbone networks in CIDR blocks (all specifics)
1849:5001        Customer networks not in CIDR blocks
1849:5005        CIDR blocks
1849:5050        Networks learned from paying peers
1849:5100        Networks learned from LINX peer ISPs
1849:5666        Multihomed customer peers
1849:6000        European peers
1849:9030        Customer networks which should only be advertised within Europe
1849:9031        Same as 9030, but 3*AS1849 prepended elsewhere
1849:9040        Customer networks which should only be advertised in the UK
1849:9041        Same as 9040, but 3*AS1849 prepended elsewhere
1849:9050        Customer networks which should only be advertised to customers
1849:9051        Same as 9050, but 3*AS1849 prepended elsewhere
```

BGP Communities

```
Community-list 1        announced to peers at regional exchange points;
                        list is made up of 1849:5001,5005 and 5006 only.
Community-list 6        forced leakage of CIDR block subnets; list contains
                        1849:5666 only
Community-list 7        set local pref 70; list contains 1849:70
Community-list 8        set local pref 80; list contains 1849:80
Community-list 9        set local pref 90; list contains 1849:90
Community-list 10       specifics originated within 1849; list contains
                        1849:5000 only
Community-list 11       set local pref 110; list contains 1849:110
Community-list 12       UK exchange point networks; list contains 1849:5100
Community-list 13       set local pref 130; list contains 1849:130
Community-list 17       all AS701 routes (no 702); list contains 1849:701
Community-list 18       all AS702 routes (no 701); list contains 1849:702
Community-list 21       the whole internet
Community-list 22       non-UK European peers; list contains 1849:6xxx
Community-list 23       routes advertised in EU only; 1849:9030
Community-list 24       as 23 but with 3*AS1849 prepend; 1849:9031
Community-list 25       routes advertised in UK only; 1849:9040
Community-list 26       as 25 but with 3*AS1849 prepend; 1849:9041
Community-list 27       routes advertised to customers only; 1849:9050
Community-list 28       as 27 but with 3*AS1849 prepend; 1849:9051
```

Sample Configurations

- **Documentation includes configurations:**
  - **border router**
  - **core router**
  - **gateway router**
- **Too much to put here!**

IP Addressing

- **UUNET UK is a European Local Internet Registry**
  - **address space delegated from RIPE**
  - **assigns address space to customers and other ISPs according to RFC2050**
- **Detailed and well documented national addressing plan**

Addressing Scheme

- **Internal Use**
  - **Point to point link address space assigned per PoP router**
  - **backbone address space assigned per region**
  - **loopback interfaces addressed from a small block of address space**
  - **security & simplicity**
- **Customer Use**
  - **assigned per need according to RFC2050**
  - **typically from /28 to /18 in size**
  - **no regional aggregation**

Services

- **Locating Servers crucial to network operation and performance**
  - **DNS**
    - **cache - for customer resolver use**
    - **primary**
    - **secondary**
  - **News**
    - **distributed to PoPs from incoming “gateway”**
    - **news peerings with major Internet sites**
    - **separate posting news server**
    - **separate farm for online news reading**

Services (continued)

- **More servers:**
  - **Authentication**
    - **RADIUS for Home User DIAL service**
    - **TACACS+ for engineering access**
  - **Mail**
    - **relay for customers**
    - **pop3 for Home User DIAL service**
  - **Logging**
    - **loghosts for all equipment**
    - **different “levels” for different systems**
- **Two of everything!**

Operations

- **ISP's need**

  organisational structure.

  operational policies.

  customer guarantees

  supplier maintenance contracts

  on-site spares

  proper test lab/environment

⇒ All part of UUNET UK's operation.

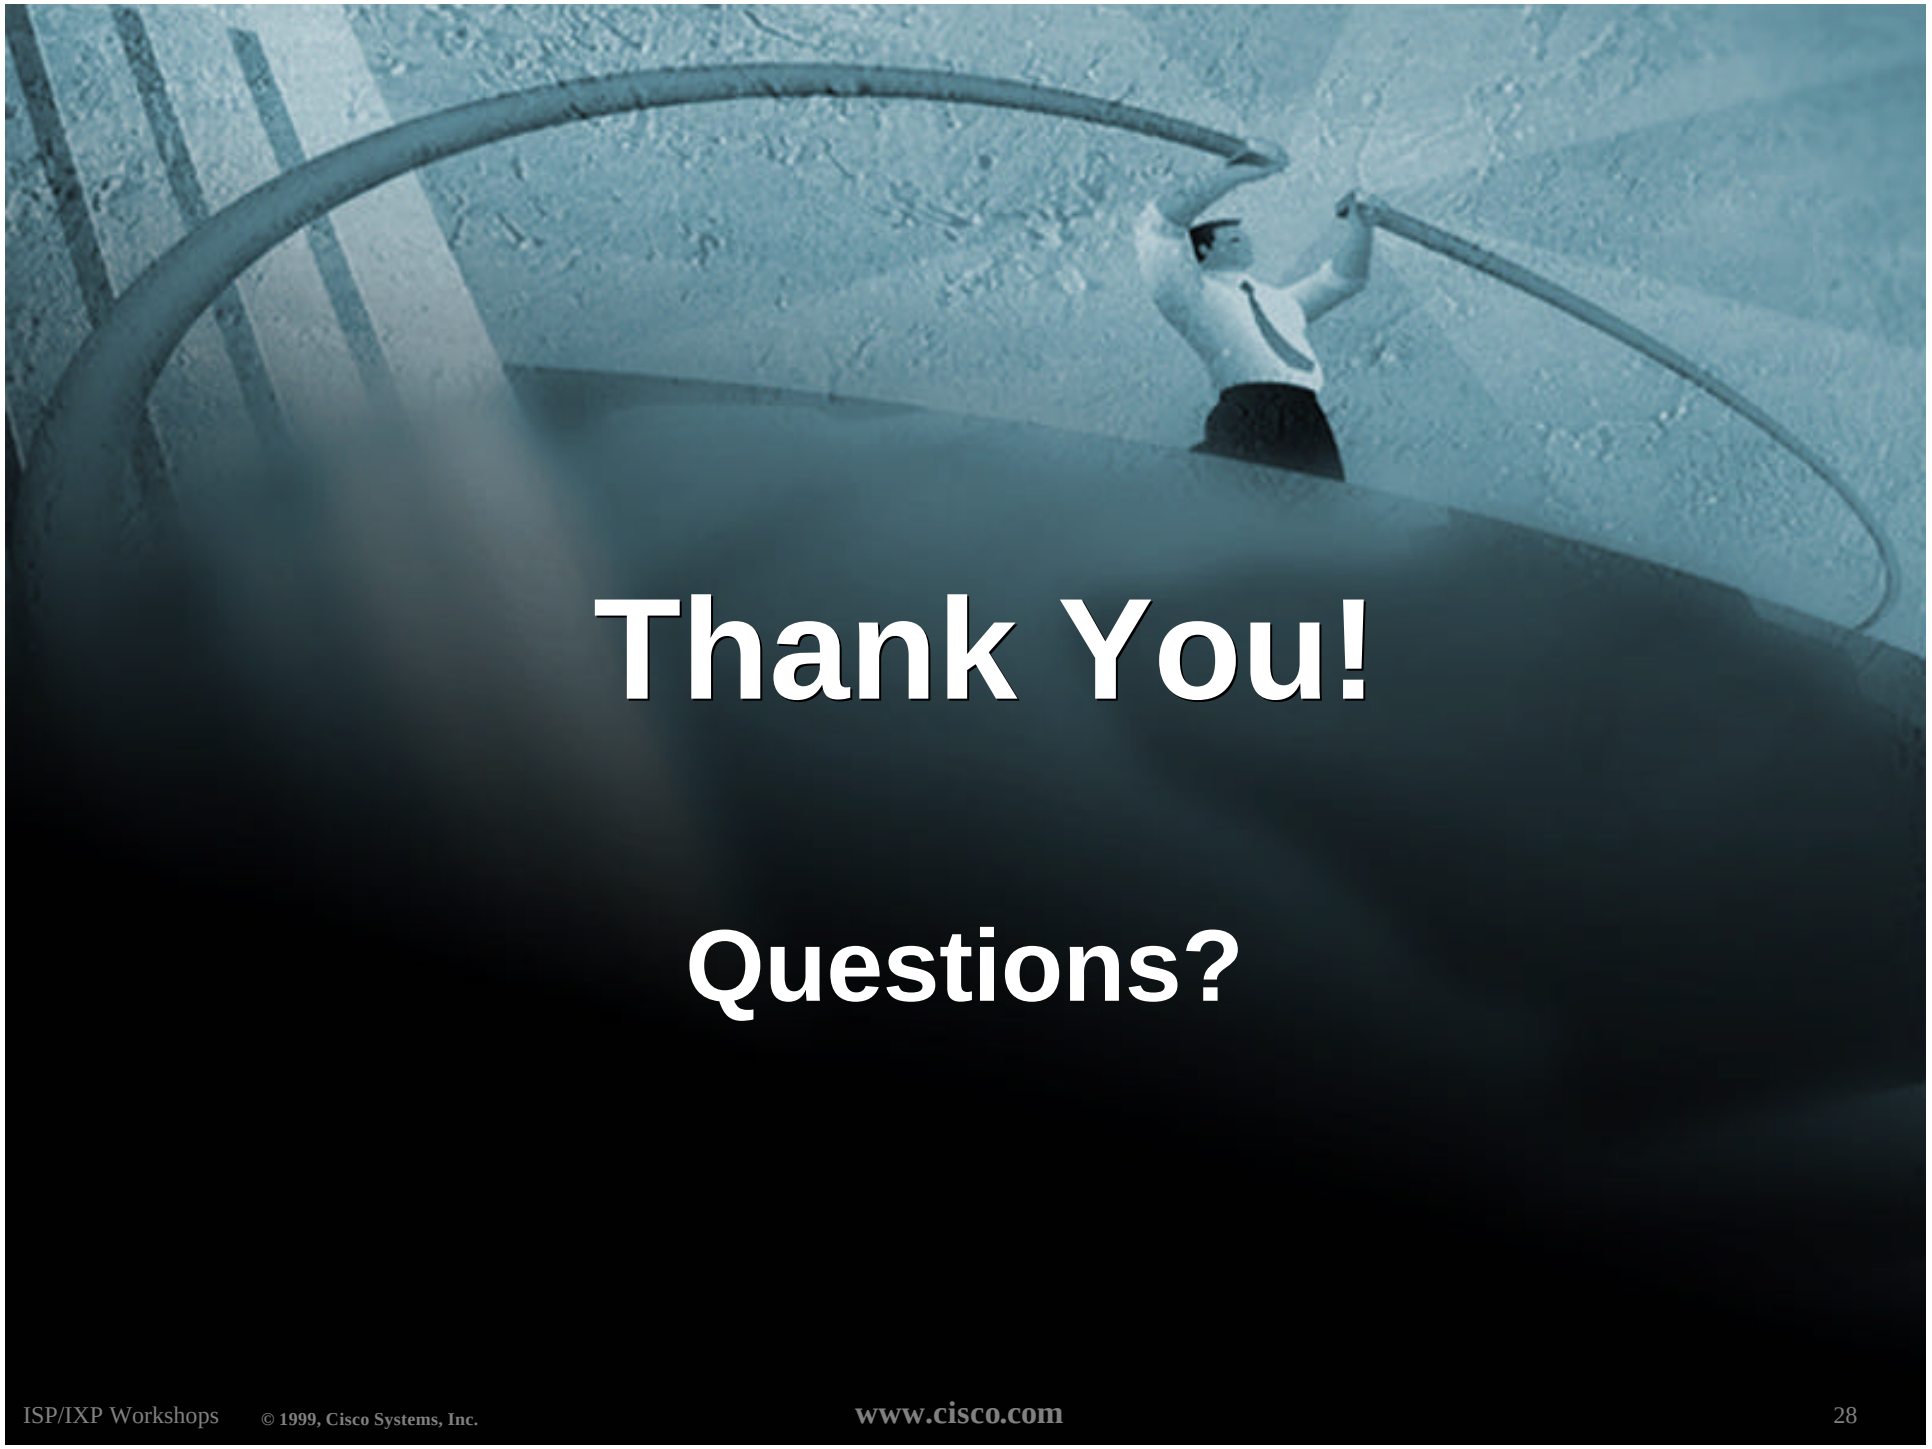

Thank You!

Questions?