ISP Case Study – PIPEX

Version 1.0

Philip Smith,
Consulting Engineering, Office of the CTO,
Cisco Systems Inc

Introduction

This document gives a brief overview into many of the design principles and configurations used at a major European ISP. The design covers much of the course work of Cisco's ISP/IXP workshop programme, and provides significant examples of some of the advanced techniques discussed during the workshop.

Thanks are due to UUNET UK[1] for permission to use these configuration examples for this work. In particular thanks are due to Judith Blair (Technical Director), and Stephen Hagger (Network Development Manager).

Network Design Principle

Background

UUNET UK is a leading UK ISP, founded early in 1992 by the software house Unipalm. Prior to purchase by UUNET, the ISP business was better known as PIPEX, and at that time was one of the major players and technology leaders in the European Internet scene.

PoPs

The network design principle was very straightforward. Keep it simple! Points of presence were located in major population centres, starting with the largest UK cities, and working into smaller locations as costs of operation decreased. The first two sites were at the HQ in Cambridge, and a major PoP in London.

PoPs were added once the financial case made sense. When there was a significant number of customers at a particular location, such that it made it more viable for PIPEX to locate equipment in that location than provide backhaul to the next closest PoP, a new PoP was

[1] UUNET UK is the current name for PIPEX (The Public IP Exchange Limited) which was the UK's first commercial ISP. PIPEX was founded in February 1992, and connected its first customer in April 1992. The author joined PIPEX in January 1993 as the support engineer – at the time, the company was a team of 7 people. Such was the success of PIPEX, that it firstly floated on the UK stockmarket with its parent company, Unipalm, and was later purchased by UUNET Technologies as part of its worldwide expansion plans. UUNET is now part of the Worldcom Group of Companies. On leaving, the author was Head of Network Engineering, having driven the company's network and product offering to be one of the most reliable and technically advanced in Europe.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

built. A couple of years ago, this was typically 60 customer 64kbps equivalent leased line circuits. More recently, with newer technology, this number is closer to 30.

Inter-PoPs

Each PoP had at least two exit routes, on diverse paths, and two other PoPs. This protected against router failure, trunk failure, and even whole PoP failure. Equipment is usually more reliable than the power supply or human "intervention", so it made sense to protect against more than just the potential equipment problems most people think of.

Rather than rely on some hazy circuit provision between sites, leased line circuits were deliberately chosen as inter-PoP links. The bandwidth on these is guaranteed, unlike X.25, Frame Relay or ATM. While costing slightly more than these 3 technologies, it meant accurate control over service quality and costs could be maintained.

Backbone

The leased line backbone was provided by as many different telephone companies as feasible. British Telecom is the major carrier in the UK, but is seeing competition from the likes of Cable and Wireless, Energis, Scottish Telecom, and the Electricity Companies. PIPEX inter-PoP backbone links would generally be provided by BT and one other. This protected the ISP against major infrastructure failure of one telco – this did happen, and not infrequently either.

Exit Points

At least two exit points to the network were provisioned. One exit point from the UK meant a single point of failure, and why build in resilience for the backbone if there was only one link to the rest of the Internet. Prior to UUNET, two external links to the US were provisioned. One was Cambridge-Washington, the other was London-New York. Apart from this, links to major European exchange points were also provisioned (the aim being to keep intra-European traffic in Europe).

PoP Design Principle

When PIPEX was first started, the business was low budget compared with many of the startups in business today. One router was purchased, and it had to carry out pretty much every function required in the network: customer aggregation, backbone links, external links. However, as the network and business grew, this model was rapidly discarded.

Core Routers

Core routers are located in the core of the network only, and carry backbone links. This is their primary function. They need to be available 100% of the time, and are optimised for higher speed circuits than any other router on the network. PIPEX's typical core router today is a 7507 with RSP4, VIP2/40 interfaces, and 128Mbytes memory. Software used is from either the CA (original ISP) or CC trains – the latter introduces Cisco Express Forwarding, a new high speed packet forwarding methodology.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

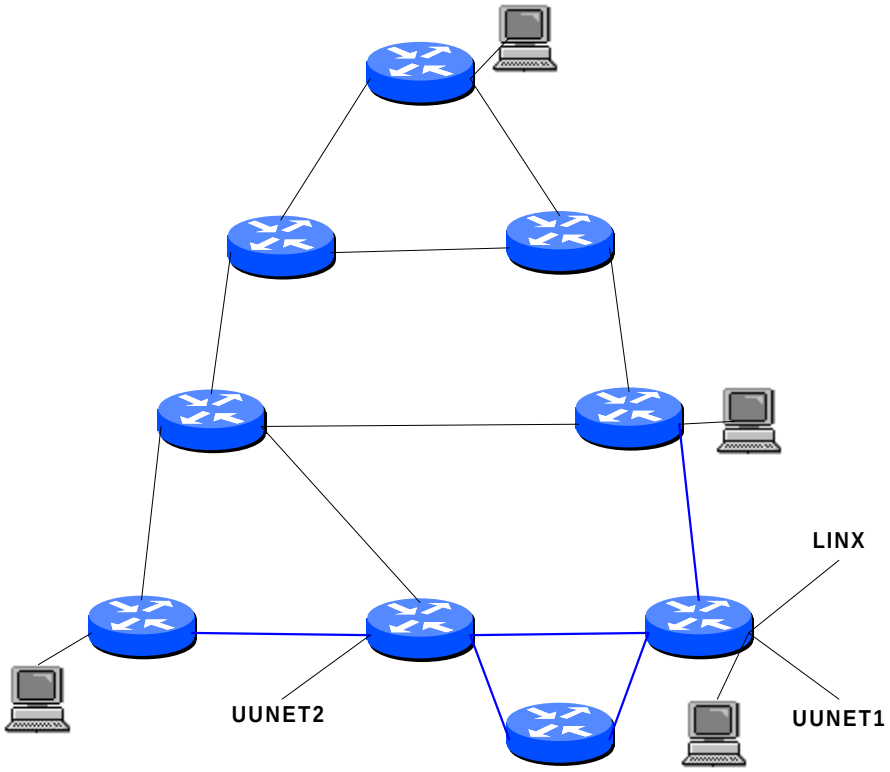

Figure 1 - national network layout

Gateway Routers

Gateway routers are located on the edge of the network only, providing the gateway between the network core and the customers' networks. These systems are optimised for the function of traffic aggregation at low to medium data rates. The typical gateway router is a 7507 with RSP2, 2x FSIP and 2x MIP interfaces, aggregating customer circuits from 64kbps up to 2Mbps. All routers run CA software, although at the time of writing plans are being drawn up to migrate to CC software.

Service Routers

Another type of aggregation router is the service router, providing the "interface" between hosted services, access and server LANs, and the core of the network. Service Routers tend to be only located in the larger PoPs where they are required. Their specification is very similar to that of the Core Router.

Border Routers

As the name suggests, these systems provide the link between the ISPs main backbone network and other external networks. These external networks may be behind transoceanic links, or at local or International exchange points. These routers are currently 7507s with RSP4, 256Mbytes RAM, and highspeed VIP2/40 interface processors. They carry large

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

numbers of peering sessions, employ policies such as route filtering, AS filtering, and route flap dampening.

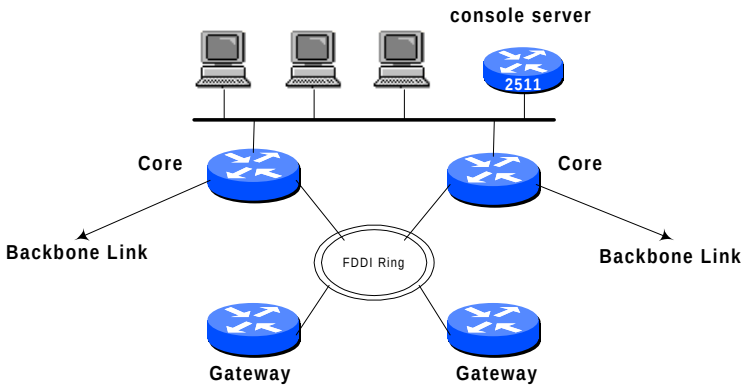

Figure 2 - small PoP

Access Routers

For the low end consumer marketplace, access routers are used to terminate PSTN (V.34 and V.90) and ISDN links. Such routers include AS5300, Ascend MAX, and 3com TC Hubs, depending on which network contract is being serviced. These routers are all “dumb” devices as far as the network is concerned, with the main routing carried out by the service routers.

PoP Design

In general, two PoP designs were employed. The first, the “small” PoP was deployed when starting a new point of presence. If the justification for siting a new PoP could pay for this equipment spend, the project would go ahead. The main inter-device connection is FDDI, although this is now being replaced by basic switched 100baseT.
The second type of PoP is the “large” PoP, typical when a site had a year or two of operation behind it. The main FDDI ring is now replaced with a pair of Catalyst 5000 switches running vLANs.

Routing Configuration

IGP Configuration

The IGP design follows the recommendations in the ISP/IXP workshop programme, and is indeed what many ISPs chose to employ today. Hierarchical!

The IGP used is OSPF. After several years running IGRP and then EIGRP, the decision was made to switch over to OSPF. The choice was mainly that OSPF is non-propietary, used by other vendors, and an IETF standard. The technology difference between distance vector (EIGRP) and link stat (OSPF) protocols did not really enter the decision process, for the reasons below.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

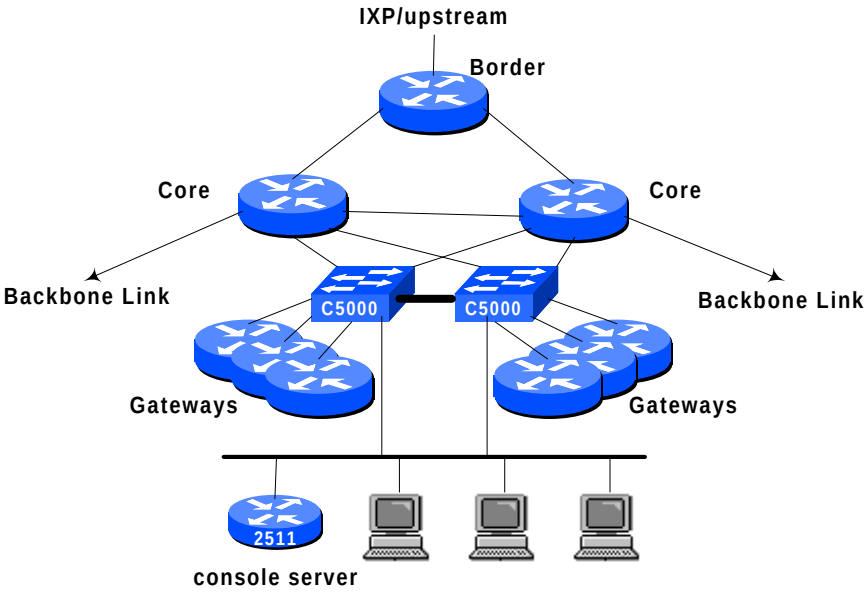

Figure 3 - Large PoP

The core backbone was defined as being in area 0. Each node in area 0 must see a path to every other node, and there cannot be any node islands. Hence the inter-core router links in the section above. No point to point link can exist in the same area at one time.

Each PoP was defined to be in a separate OSPF area. Hierarchical structure. Networks are summarised at area boundaries, resulting in very few networks floating around in area0. This gives rapid convergence in case of link failure, so much so that link failures are barely noticed today.

The same was done with EIGRP. Networks used in each PoP were summarised on the links between PoPs. The core EIGRP was kept small, also giving rapid convergence. The argument between the merits of link state versus distance vector? A draw, for PIPEX.

BGP Configuration

Route Reflectors

Similar attention was applied to the BGP configuration. Each PoP was its own route reflector cluster, with the core routers being the reflectors, the other routers being clients. In a PoP the iBGP was fully meshed, but for the backbone only the route reflectors are fully meshed. Rather than having 100 routers all partaking in the iBGP, there are 10 or more clusters of only 10 routers each. Hierarchical and scalable!

Further, this design makes it easy to add routers to the network. The iBGP for the whole network doesn't need to be reconfigured, only the particular PoP in question. Same for adding a new PoP.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

Full Internet Routes

Only the core routers in the PoP carry full Internet Routing. The other routers only see UK routes. There is indeed no need for the core router to see full Internet Routes, but some customer pay for a multihoming service which includes that functionality. Also, there is no need for the non-core routers to carry UK routes, but that was a design choice.

Border Routers

Great care is required when configuring BGP for border routers. Indeed, the design chosen here was that the border routers which connected to the upstream ISP (transit provider), ie UUNET Europe, were **separate** devices from those which connected to regional and local ISPs using either private or public peering points. Only the border router which connected to the transit provider carried full internet routers – the routers connecting to the exchange point or private peers carried only UUNET UK routes and the routes they heard from the external neighbours. This avoided the local peering routers from being used as default gateways by other domestic or regional ISPs. (There were several recorded instances of this, hence the design decision.) Furthermore, the local border routers did not carry a default route.

Default Routes

As UUNET UK carried full Internet routes throughout the core of the network, there was no need to have a default route configured. However, the decision was made to point the default to the transit ISP (UUNET Europe) as many customers didn't understand, nor could be persuaded, that a non-announced route in the Internet wasn't our fault. Having their traceroute terminate in "someone elses network" was acceptable to them, though! Also, border routers connecting to private peers or local/regional exchange points have no default route configured, avoiding potential traffic dumping.

BGP Route Flap Dampening

Route flap dampening as per RIPE-178 is applied on the network border routers. This minimises the ripple effect of the constant changing in the Internet Routing table from affecting the whole of the domestic network.

The full configuration used for route flap dampening is given below:

basic bgp configuration and implementation of route-map

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

no flap dampening for special 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 (30 mins half life, 750 reuse, 3000 penalty, reuse 60 mins)

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

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

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

```
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 will be applied using the default dampening configuration.

BGP Communities

Extensive use of BGP communities is made for applying different tags to different types of customer networks. The following list is the range of different communities which were deemed necessary:

```
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
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

Further, it is possible to bundle different communities together into community lists – these are defined below. Community lists are processed at peerings with other AS's, to force particular actions for particular requirements.

```
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
```

Address and Routing Registries

UUNET UK is a local Internet Registry, delegated by RIPE to assign address space to its customers in the UK. Similar models exist for the other Regional Registries (APNIC, ARIN).

Networks UUNET UK assign to customers were stored in a local database, and stored in the RIPE address registry. This is essential for documentation purposes so that Internet users know who has what address space in case of technical problems, connectivity issues, and so forth.

Separate from the address registry is the Internet Routing Registry. This is a global database listing Routing Policy for various networks. UUNET UK makes extensive use of the IRR, specifically the registry operated by the RIPE NCC on behalf of the European ISP community. Route objects and AS objects for UUNET UK, and its customers are registered there. Many ISPs filter on what is registered in the Internet Routing Registry, hence the importance of keeping the information accurate and uptodate. Often a failure to register meant a failure for the affected customer from getting connectivity to many parts of the world. The IRR is also a very important tool, used by UUNET UK's operations team to debug many network problems on behalf of their customers.

Services Locations

While routing configuration is key to a successful ISP business, some thought needs to be applied to location of key services. But what is meant by services? Exactly the same principles apply to locating servers as does to locating routers. The minimum of two of anything ensures reliability and service continuity in case of faults. Placing key servers, such as DNS, very close to the core of the network

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

ensures no delay in response to queries. And distributing servers such as those providing newsfeeds so that there is minimused utilisation of long distance bandwidth provides great advantage.

DNS Servers

DNS RFCs state and best common practice by most ISPs is that any domain name has at least two name servers for it. One name server should be on the local ISP's network, the other name server should be with an upstream provider for resilience purposes. Several ISPs offer a service of "mirroring" primary nameservers, and indeed most large ISPs will offer this service to their customers (UUNET is one example).

However, should an ISP only have one name server on its local network. UUNET UK chose not to do this as it is very unsafe, doesn't give reliable service, and simply does not scale. At the end of 1997, UUNET UK was hosting 26000 domains, a considerable number, hence requiring a bit of thought to deploy and operate reliably.

The current DNS software is very flexible and allows for a scalable DNS solution. The primary nameservers at UUNET UK are kept away from customers and the general public view (some ISPs choose to place them behind a firewall), and only reply to requests from "caching nameservers" which are the public front end. Nameserver caches are fast computer systems with large amounts of memory, able to provide DNS responses very rapidly.

Authentication Servers

Not simply one computer siting in a corner, but a network of computers providing authorisation, authentication, and accounting for user connections.

Backbone Routers

Three TACACS+ servers provided authentication for the backbone routers. Two were "public", with a master server sitting in a protected network feeding the slaves. These servers authenticated only the users connecting to backbone routers – no public user information was kept on these, for obvious security reasons.

Services

Each dialup service had its own RADIUS server. Each dialup network had its own RADIUS proxy, which knew how to get to the RADIUS server for the service in question. This allowed UUNET UK to offer MSN, AlterDial, PIPEX DIAL, and other roaming services on its network. Note the use of proxy servers, backended by separate RADIUS servers. This scaled much better than trying to use one or two systems. Accounting information was sent to a system separate from the authentication servers – again it is better to separate functions out between multiple different systems than trying to run everything on one.

Mail

Likewise the mail system functionality was spread between many systems rather than trying to do everything on one. Dialup services POP3 server was separate from the SMTP system used by users to send e-mail. Mail relay services for customers (backup MX records) were separate from mailsystems which handled company mail. Company public mail was separated from company internal mail. All this was implemented to spread the load amongst different systems, improve security, improve reliablity.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

Note that one MX record for a mailhost is accepted bad practice. At least two were configured for all customers, one being the customer system, the other being the local relay. Often more than two were configured at the customer's request.

News

The design of the news (NNTP) systems was one borne out of keeping up with the tidal wave. The current design was to install at least one news system in each PoP, adding more as needed. This avoided sending multiple large newsfeeds over the backbone, offered better resilience, and scaled. External newsfeeds came into a single collector, which then distributed the feed out to each PoP. News postings or incoming feeds from customers came to a different system optimised to receive feeds. News machines were also provided for dialup users who wanted a service allowing them to read news – again a cluster of separate systems.

Operations

Operating a network is considerably beyond the scope of this document. However, some pointers are worth noting. UUNET UK split into several business functions once it reached around 15 staff – the company no longer scaled with everyone doing a bit of everything.

Every ISP needs an organisational structure once they are properly established. A development engineer can't spend all her time operating the network, as a customer support engineer can't spend all his time selling services. Structure – split technical and sales operations, and subdivisions within there allows everyone to do their job to the full.

Within the technical side, larger companies need proper operating policies. It is no good an operations engineer replacing a piece of hardware while the network is live, or if the hardware concerned is going to be replaced by someone else in a larger programme of work in a few weeks. Agree on regular maintenance slots (4am to 7am in the local timezone is a good time for minimal customer disruption), implement a change control system so that everyone working on the network documents what they are changing and the reasons why. Ensure that network changes are reviewed by key engineering staff – avoids duplication of effort, reduces chance of error.

Offer customers a service guarantee. UUNET UK offered 99.5% guaranteed availability within the network they controlled (upto and including the customer's router if it was supplied by UUNET UK) – while this is a key service differentiator, it backs up the claims and engineering effort put into making a reliable network. UUNET UK never had total outage which dropped below the above guarantee.

Ensure that proper maintenance contracts are set up with key suppliers. When the key router breaks late Friday night, no ISP wants to wait until Monday morning to log a call under a cheap 9am-5pm contract. And having a maintenance contract is not enough – UUNET UK kept local spares at each PoP, and in a central location, just in case things went badly wrong. The times they did, service could be returned within a couple of hours, not waiting for the 4 hour guaranteed response a maintenance contract offered (a response could simply mean acknowledgement of e-mail receipt).

Set up a test lab environment, ideally a duplicate of a typical PoP. UUNET UK implemented this as part of its sparing policy. In the event of hardware failure, the development PoP could be raided to

supply essential spares. And new code could be tested in complete isolation from the main network before it was connected to the main network, or even deployed. Very important to avoid catastrophic network failure, or even denial of service to customers due to some undetected software or hardware bug.

All the above were utilised to the full by UUNET UK, in addition to stringent network design. The package as a whole meant a successful ISP, offering levels of service acceptable to most of its customers, and almost unparalleled by other “cheaper” ISPs in the UK.

Summary

This document has taken a brief look at some of the design principles and configurations used on the UUNET UK network. Routing and systems configurations have been discussed, and while this document is by no means a complete description of the network, it is hoped that a flavour of what is involved in configuring a reliable, resilient, and robust network has been given.

No detailed description has been given of organisational structure, or operational policies. These are also key to the operation of a successful ISP, but are considerably beyond the scope of this technical document.

Questions and requests for further clarification can be direct to the author – e-mail pfs@cisco.com.

Further Reading

1. RFC1918
2. RFC1998
3. IOS Essentials for ISPs – Cisco ISP/IXP Workshop programme
4. Internet Routing Architectures – Bassam Halabi – Cisco Press, ISBN 1-56205-652-2.

Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

Appendix – Sample Configurations

Border Router

This is a sample configuration from one of the border routers. Irrelevant or sensitive information has been removed or altered. Comments are in **bold** font.

Handy comment of when the changes were made, and by whom.

```
!
! Last configuration change at 08:16:49 BST Thu Apr 23 1998 by <removed>
! NVRAM config last updated at 08:17:30 BST Thu Apr 23 1998 by <removed>
!
```

Essential changes for an ISP

```
version 11.1
no service finger
service timestamps debug datetime
service timestamps log datetime
service password-encryption
no service udp-small-servers
no service tcp-small-servers
!
hostname doc-br1
!
clock timezone GMT 0
clock summer-time BST recurring last Sun Mar 1:00 last Sun Oct 1:00
!
```

Set up authentication, authorisation, and accounting – tacacs+

```
aaa new-model
aaa authentication login default tacacs+ enable
aaa authentication enable default tacacs+ enable
aaa accounting exec start-stop tacacs+
aaa accounting commands 15 start-stop tacacs+
enable secret 5 <removed>
enable password 7 <removed>
!
```

More essential changes for an ISP

```
ip subnet-zero
ip spd enable
ip ftp source-interface Loopback0
ip ftp username <removed
ip ftp password 7 <removed>
ip tacacs source-interface Loopback0
!
```

Loopback interfaces never go away!

```
interface Loopback0
 ip address 158.43.206.96 255.255.255.255
!
```

Interface to the Core routers

```
interface Fddi0/0/0
 description Border FDDI ring for doc-cr1/2 and doc-br1/2
 ip address 158.43.195.3 255.255.255.0
 no ip directed-broadcast
 no ip route-cache optimum
 ip route-cache flow
 ip route-cache distributed
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
 no keepalive
!
```

Interface to the upstream ISP (UUNET Europe)

```
interface Fddi0/1/0
 description FDDI to UUNET Superhub
 ip address 146.188.31.200 255.255.255.224
 ip access-group 199 in                        ! block smurf
 no ip directed-broadcast
 no ip route-cache optimum
 ip route-cache flow
 ip route-cache distributed
 no keepalive
!
```

Interface to the London Exchange Point

```
interface Fddi1/0/0
 description FDDI to linx-br2
 ip address 158.43.194.129 255.255.255.128
 ip access-group 199 in
 no ip directed-broadcast
 no ip route-cache optimum
 ip route-cache flow
 ip route-cache distributed
 ip ospf cost 100
 delay 400
 no keepalive
!
```

Interface to the Microsoft European DataCentre

```
interface FastEthernet4/0/0
 description 100bFX link to Microsoft DataCentre
 ip address 193.128.43.5 255.255.255.252
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 no ip route-cache optimum
 ip route-cache flow
 ip route-cache distributed
 media-type MII
 full-duplex
!
```

For information – our AS!

```
autonomous-system 1849
!
```

OSPF configuration

```
router ospf 44
 redistribute connected subnets route-map connected-to-ospf
 redistribute static subnets route-map static-to-ospf
 passive-interface Fddi0/1/0
 passive-interface FastEthernet1/1/0
 passive-interface FastEthernet4/0/0
 passive-interface Loopback0
 network 158.43.192.0 0.0.15.255 area 20
 network 158.43.64.0 0.0.15.255 area 20
 maximum-paths 6
 default-metric 20
!
```

BGP configuration

```
router bgp 1849
!
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

IOS Essentials

```
no synchronization
no bgp fast-external-fallover
bgp dampening route-map expo-flap-dampen
!
```

redistribute static into BGP

```
redistribute static route-map static-bgp
```

partial/UK routing with community usage

```
neighbor core-ibgp-partial peer-group
neighbor core-ibgp-partial remote-as 1849
neighbor core-ibgp-partial update-source Loopback0
neighbor core-ibgp-partial send-community
neighbor core-ibgp-partial route-map core-ibgp-partial-out out
```

full routing with community usage

```
neighbor core-ibgp-full peer-group
neighbor core-ibgp-full remote-as 1849
neighbor core-ibgp-full update-source Loopback0
neighbor core-ibgp-full send-community
```

dual peering with AS702, UUNET Europe

```
neighbor 146.188.31.193 remote-as 702
neighbor 146.188.31.193 send-community
neighbor 146.188.31.193 distribute-list 150 in
neighbor 146.188.31.193 distribute-list 168 out
neighbor 146.188.31.193 route-map uunet-peer-in in
neighbor 146.188.31.193 route-map uunet-peer-out out
neighbor 146.188.31.194 remote-as 702
neighbor 146.188.31.194 send-community
neighbor 146.188.31.194 distribute-list 150 in
neighbor 146.188.31.194 distribute-list 168 out
neighbor 146.188.31.194 route-map uunet-peer-in in
neighbor 146.188.31.194 route-map uunet-peer-out out
```

peering with domestic core network

```
neighbor 158.43.131.104 peer-group core-ibgp-full
neighbor 158.43.131.105 peer-group core-ibgp-full
neighbor 158.43.162.104 peer-group core-ibgp-full
neighbor 158.43.162.105 peer-group core-ibgp-full
neighbor 158.43.179.104 peer-group core-ibgp-full
neighbor 158.43.179.105 peer-group core-ibgp-full
neighbor 158.43.206.97  peer-group core-ibgp-full
neighbor 158.43.206.103 peer-group core-ibgp-partial
neighbor 158.43.206.104 peer-group core-ibgp-full
neighbor 158.43.206.105 peer-group core-ibgp-full
neighbor 158.43.211.104 peer-group core-ibgp-full
neighbor 158.43.211.105 peer-group core-ibgp-full
neighbor 158.43.219.104 peer-group core-ibgp-full
neighbor 158.43.219.105 peer-group core-ibgp-full
neighbor 158.43.227.104 peer-group core-ibgp-full
neighbor 158.43.227.105 peer-group core-ibgp-full
neighbor 158.43.234.96  peer-group core-ibgp-full
neighbor 158.43.234.104 peer-group core-ibgp-full
neighbor 158.43.234.105 peer-group core-ibgp-full
neighbor 158.43.239.106 peer-group core-ibgp-full
neighbor 158.43.251.28  peer-group core-ibgp-full
neighbor 158.43.251.29  peer-group core-ibgp-full
```

peering with Microsoft European DataCentre

```
neighbor 193.128.43.6 remote-as 8068
neighbor 193.128.43.6 soft-reconfiguration inbound
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
  neighbor 193.128.43.6 distribute-list 101 in
  neighbor 193.128.43.6 distribute-list 168 out
  neighbor 193.128.43.6 route-map microsoft-dc-in in
  neighbor 193.128.43.6 route-map uk-transit-out out
```

peering with RouteViews at University of Oregon

```
  neighbor 198.32.162.100 remote-as 65534
  neighbor 198.32.162.100 ebgp-multihop 255
  neighbor 198.32.162.100 update-source Loopback0
  neighbor 198.32.162.100 next-hop-self
  neighbor 198.32.162.100 distribute-list 150 in
  neighbor 198.32.162.100 distribute-list 168 out
  neighbor 198.32.162.100 route-map uni-oregon-in in
  neighbor 198.32.162.100 route-map full-routing-out out
```

miscellaneous BGP commands – note modified distances

```
  maximum-paths 2
  distance bgp 180 200 200
  no auto-summary
 !
```

configuration helpers

```
 ip host NAME1 a.b.c.d
 ip host NAME2 v.w.x.y
 ip domain-name pipex.net
 ip name-server 158.43.128.1
 ip name-server 158.43.192.1
 ip name-server 198.6.1.1
```

NetFlow configuration

```
 ip flow-export destination a.b.c.d port
 ip flow-export source Loopback0
 ip flow-export version 5 origin-as
```

IOS Essentials – CIDR and Martians

```
 ip classless
 ip route 1.0.0.0 255.0.0.0 Null0
 ip route 10.0.0.0 255.0.0.0 Null0
 ip route 19.255.0.0 255.255.0.0 Null0
 ip route 59.0.0.0 255.0.0.0 Null0
 ip route 89.0.0.0 255.0.0.0 Null0
 ip route 99.0.0.0 255.0.0.0 Null0
 ip route 125.0.0.0 255.0.0.0 Null0
 ip route 127.0.0.0 255.0.0.0 Null0
 ip route 129.156.0.0 255.255.0.0 Null0
 ip route 172.16.0.0 255.240.0.0 Null0
 ip route 192.5.0.0 255.255.255.0 Null0
 ip route 192.9.99.0 255.255.255.0 Null0
 ip route 192.9.200.0 255.255.255.0 Null0
 ip route 192.168.0.0 255.255.0.0 Null0
 ip route 223.255.255.0 255.255.255.0 Null0
```

Definition of community lists

```
 ip bgp-community new-format
 ip community-list 1 permit 1849:5001
 ip community-list 1 permit 1849:5005
 ip community-list 1 permit 1849:5666
 ip community-list 1 deny internet
 ip community-list 6 permit 1849:5666
 ip community-list 7 permit 1849:70
 ip community-list 8 permit 1849:80
 ip community-list 9 permit 1849:90
 ip community-list 10 permit 1849:5000
 ip community-list 11 permit 1849:110
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
ip community-list 12 permit 1849:5100
ip community-list 13 permit 1849:130
ip community-list 17 permit 1849:701
ip community-list 18 permit 1849:702
ip community-list 21 permit internet
ip community-list 22 permit 1849:6000
ip community-list 22 permit 1849:6100
ip community-list 22 permit 1849:6200
ip community-list 22 permit 1849:6300
ip community-list 22 permit 1849:6400
ip community-list 22 permit 1849:6500
ip community-list 22 permit 1849:6600
ip community-list 22 permit 1849:6700
ip community-list 22 permit 1849:6800
ip community-list 22 permit 1849:6900
ip community-list 23 permit 1849:9030
ip community-list 24 permit 1849:9031
ip community-list 25 permit 1849:9040
ip community-list 26 permit 1849:9041
ip community-list 27 permit 1849:9050
ip community-list 28 permit 1849:9051
ip community-list 28 permit 1849:9055
```

Definition of AS_PATH access-lists

- used for same AS multihomed customers

```
ip as-path access-list 2 permit ^(2830_)+$
ip as-path access-list 2 deny .*
```

- used for defining which networks are from UUNET Europe only

```
ip as-path access-list 23 deny ^702_701_
ip as-path access-list 23 deny ^702_703_
ip as-path access-list 23 deny ^702_704_
ip as-path access-list 23 deny ^702_705_
ip as-path access-list 23 permit ^702_
```

- default deny

```
ip as-path access-list 99 deny .*
!
```

IOS Essentials

```
ip ospf name-lookup
logging buffered 65536
logging trap debugging
logging source-interface Loopback0
logging <removed>
logging <removed>
```

Standard Access List definitions

- used by OSPF and others – 158.43/16 is PIPEX net for backbone point to point links

```
access-list 2 permit 158.43.0.0 0.0.255.255
access-list 2 deny   any
```

- vty security

```
access-list 3 permit <removed>
access-list 3 deny   any
```

- “safe” networks used in OSPF redistribution

```
access-list 6 permit 146.188.0.0 0.0.255.255
access-list 6 permit 193.128.40.0 0.0.7.255
access-list 6 permit 194.68.130.0 0.0.1.255
access-list 6 deny   any
```

- “safe” networks used in external peerings

```
access-list 7 permit 194.68.130.0
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
access-list 7 permit 146.188.0.0
access-list 7 permit 137.39.0.0
access-list 7 permit 158.43.0.0
access-list 7 deny   any
```

- SNMP access definition

```
access-list 98 permit <removed>
access-list 98 deny   any
```

- Default Deny!

```
access-list 99 deny   any
!
```

Extended Access Lists Defintion

- nets acceptable from Microsoft European DataCentre

```
access-list 101 permit ip host 207.46.0.0 host 255.255.224.0
access-list 101 permit ip host 207.46.32.0 host 255.255.224.0
access-list 101 deny   ip any any
```

- Access List defining PIPEX net blocks

```
access-list 143 permit ip host 158.43.0.0 host 255.255.0.0
access-list 143 permit ip host 193.128.0.0 host 255.252.0.0
access-list 143 permit ip host 193.132.0.0 host 255.254.0.0
access-list 143 permit ip host 194.128.0.0 host 255.252.0.0
access-list 143 permit ip host 194.200.0.0 host 255.252.0.0
access-list 143 permit ip host 194.216.0.0 host 255.255.0.0
access-list 143 permit ip host 195.217.0.0 host 255.255.0.0
access-list 143 deny   ip any any
```

- Access List defining *specifics* of PIPEX net blocks

```
access-list 144 permit ip 193.128.0.0 0.3.255.255 255.252.0.0 0.3.255.255
access-list 144 permit ip 193.132.0.0 0.1.255.255 255.254.0.0 0.1.255.255
access-list 144 permit ip 194.128.0.0 0.3.255.255 255.252.0.0 0.3.255.255
access-list 144 permit ip 194.200.0.0 0.3.255.255 255.252.0.0 0.3.255.255
access-list 144 permit ip 194.216.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 144 permit ip 195.217.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 144 deny   ip any any
```

- Access List defining specifics of backbone, RFC1918 and Martian networks

```
access-list 145 deny   ip host 0.0.0.0 any
access-list 145 deny   ip 0.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 145 deny   ip 1.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 145 deny   ip 10.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 145 deny   ip 59.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 145 deny   ip 89.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 145 deny   ip 99.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 145 deny   ip 125.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 145 deny   ip 127.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 145 deny   ip 129.156.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 145 deny   ip 158.43.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 145 deny   ip 172.16.0.0 0.15.255.255 255.240.0.0 0.15.255.255
access-list 145 deny   ip 192.5.0.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 145 deny   ip 192.9.200.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 145 deny   ip 192.9.99.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 145 deny   ip 192.168.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 145 deny   ip 208.205.11.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 145 permit ip any any
```

- RFC1918 and backbone nets – used in inbound peerings

```
access-list 150 deny   ip host 0.0.0.0 any
access-list 150 deny   ip any 255.255.255.128 0.0.0.127
access-list 150 deny   ip 10.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 150 deny   ip 172.16.0.0 0.15.255.255 255.240.0.0 0.15.255.255
access-list 150 deny   ip 192.168.0.0 0.0.255.255 255.255.0.0 0.0.255.255
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
access-list 150 deny   ip 158.43.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 150 permit ip any any
```

- Martian, RFC1918 and other bad nets – used in outbound peerings

```
access-list 168 deny   ip host 0.0.0.0 any
access-list 168 deny   ip 0.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 168 deny   ip 1.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 168 deny   ip 10.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 168 deny   ip 19.255.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 168 deny   ip 59.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 168 deny   ip 127.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 168 deny   ip 128.0.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 168 deny   ip 129.156.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 168 deny   ip 172.16.0.0 0.0.15.255 255.255.240.0 0.0.15.255
access-list 168 deny   ip 191.255.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 168 deny   ip 192.0.2.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 168 deny   ip 192.5.0.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 168 deny   ip 192.9.200.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 168 deny   ip 192.9.99.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 168 deny   ip 192.168.0.0 0.0.255.255 255.255.0.0 0.0.255.255
access-list 168 deny   ip 223.255.255.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 168 deny   ip 64.0.0.0 31.255.255.255 224.0.0.0 31.255.255.255
access-list 168 deny   ip 96.0.0.0 31.255.255.255 224.0.0.0 31.255.255.255
access-list 168 deny   ip any 255.255.255.128 0.0.0.127
access-list 168 permit ip any any
```

- Access Lists for route flap dampening as per RIPE-178 definition

```
access-list 180 permit ip host 198.41.0.0 host 255.255.252.0
access-list 180 permit ip host 128.9.0.0 host 255.255.0.0
access-list 180 permit ip host 192.33.4.0 host 255.255.255.0
access-list 180 permit ip host 128.8.0.0 host 255.255.0.0
access-list 180 permit ip host 192.203.230.0 host 255.255.255.0
access-list 180 permit ip host 192.5.4.0 host 255.255.254.0
access-list 180 permit ip host 192.112.36.0 host 255.255.255.0
access-list 180 permit ip host 128.63.0.0 host 255.255.0.0
access-list 180 permit ip host 192.36.148.0 host 255.255.255.0
access-list 180 permit ip host 193.0.14.0 host 255.255.255.0
access-list 180 permit ip 198.32.64.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 180 permit ip 198.32.65.0 0.0.0.255 255.255.255.0 0.0.0.255
access-list 181 permit ip any 255.255.255.0 0.0.0.255
access-list 181 deny   ip any any
access-list 182 permit ip any 255.255.252.0 0.0.3.255
access-list 182 deny   ip any any
access-list 183 permit ip any 255.255.240.0 0.0.15.255
access-list 183 deny   ip any any
access-list 188 permit ip 192.0.0.0 0.255.255.255 255.0.0.0 0.255.255.255
access-list 188 deny   ip any any
access-list 189 permit ip host 137.39.0.0 host 255.255.0.0
access-list 189 permit ip host 146.188.0.0 255.255.0.0 0.0.255.255
access-list 189 permit ip host 158.43.0.0 host 255.255.0.0
access-list 189 permit ip host 194.68.128.0 host 255.255.0.0
access-list 189 permit ip 194.68.130.0 0.0.1.0 255.255.0.0 0.0.1.0
access-list 189 deny   ip any any
```

Block SMURF attacks

```
access-list 199 deny   ip any 0.0.0.255 255.255.255.0 log
access-list 199 deny   ip any 0.0.0.0 255.255.255.0 log
access-list 199 permit ip any any
```

TACACS+ definition

```
tacacs-server host <removed>
tacacs-server host <removed>
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
tacacs-server key <removed>
!
```

Route Map defintions

- full routing to peer: aggregates plus full Internet routes

```
route-map full-routing-out deny 10
 match community 10 exact-match
!
route-map full-routing-out permit 30
 match community 21
!
```

- partial routes for iBGP peers

```
route-map core-ibgp-partial-out permit 10
 match community 1 10 12 18 22
!
```

- redistribute static routes to ospf:

```
route-map static-to-ospf permit 10
 match ip address 2
!
```

- redistribute connected networks to ospf:

```
route-map connected-to-ospf permit 10
 match ip address 6
!
```

- redistribute static routes into BGP:

```
route-map static-bgp permit 5
 match ip address 143
 set origin igp
 set community 1849:5005
!
route-map static-bgp permit 10
 match ip address 144
 set origin igp
 set community 1849:5000 no-export
!
route-map static-bgp permit 20
 match ip address 145
 set origin igp
 set community 1849:5001
!
```

- networks to announce to UUNET Europe upstream.

```
route-map uunet-peer-out deny 10
 match community 25 27
!
route-map uunet-peer-out permit 20
 match community 23
 set community 702:30
!
route-map uunet-peer-out permit 30
 match community 24
 set community 702:3
!
route-map uunet-peer-out permit 40
 match community 26 28           ! path length stuffing for those want it
 set as-path prepend 1849 1849 1849
!
route-map uunet-peer-out permit 50
 match community 1
 set metric 5
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
!
```

- networks learned from UUNET Europe upstream

```
route-map uunet-peer-in permit 5
 match ip address 7
 set community 1849:702 1849:5000
!
route-map uunet-peer-in permit 10
 match as-path 23
 set community 1849:702
!
route-map uunet-peer-in permit 20
 set community 1849:701
!
```

- networks announced to customers paying for UK transit only

```
route-map uk-transit-out deny 5
 match ip address 144
 match as-path 2
!
route-map uk-transit-out permit 10
 match community 1 12
!
```

- networks learned from RouteViews

```
route-map uni-oregon-in permit 10
 set weight 65535
 set community no-advertise
!
```

- networks learned from Microsoft European DataCentre

```
route-map microsoft-dc-in permit 10
 set community 1849:5666 1849:9040
!
```

- route flap dampening as per RIPE-178 definition

```
route-map expo-flap-dampen deny 5
 match ip address 189
!
route-map expo-flap-dampen deny 7
 match ip address 180
!
route-map expo-flap-dampen permit 9
 match ip address 188
 set dampening 30 750 3000 60
!
route-map expo-flap-dampen permit 10
 match ip address 181
 set dampening 30 750 3000 60
!
route-map expo-flap-dampen permit 20
 match ip address 182
 set dampening 15 750 3000 45
!
route-map expo-flap-dampen permit 40
 set dampening 10 1500 3000 30
!
```

Careful with the SNMP configuration

```
snmp-server community <removed> RO 98
snmp-server trap-source Loopback0
snmp-server trap-authentication
snmp-server host a.b.c.d <removed>
snmp-server host e.f.g.h <removed>
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
snmp-server host i.j.k.l <removed>
!
```

IOS Essentials

```
banner login _

Authorised access only

This system is the property of UUNET UK

Disconnect IMMEDIATELY if you are not an authorised user !

Contact noc@uk.uu.net +44 541 588638 for help.

_
!
```

Terminal line configuration

```
line con 0
 exec-timeout 3 0
 transport preferred none
line aux 0
 transport input all
line vty 0 4
 access-class 3 in
 exec-timeout 30 0
 transport preferred none
!
```

Rest of Core Dump configuration

```
exception protocol ftp
exception dump a.b.c.d
!
```

Time Synchronisation

```
ntp clock-period 17179715
ntp update-calendar
ntp peer 158.43.128.33
ntp peer 158.43.128.66
ntp peer 158.43.192.66
end
```

Core Router

This section looks at a typical core router configuration – this router is connected to the border router discussed above by FDDI. Configuration which has already been discussed above as been omitted for the sake of brevity.

```
!
hostname doc-cr1
!
```

Multicast enabled

```
ip multicast-routing
ip dvmrp route-limit 7000
!
```

Loopbacks never go away!

```
interface Loopback0
 ip address 158.43.206.104 255.255.255.255
 transmit-buffers backing-store
!
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

Connection to the Gateway Routers in the PoP

```
interface FastEthernet0/0/0
 description Docklands PoP Core Fast Ethernet
 ip address 158.43.200.1 255.255.255.0
 no ip redirects
 no ip directed-broadcast
 ip route-cache flow
 no ip route-cache optimum
 ip route-cache distributed
 full-duplex
!
```

Connection to the DNS/Mail/News servers in the PoP

```
interface FastEthernet0/1/0
 description PIPEX server Ethernet Backbone (C5000)
 ip address 158.43.193.126 255.255.255.192 secondary
 ip address 158.43.192.192 255.255.255.0
 no ip redirects
 no ip directed-broadcast
 ip route-cache distributed
 delay 15
 full-duplex
```

- HSRP configuration

```
 standby 10 priority 150
 standby 10 preempt
 standby 10 ip 158.43.192.62
 standby 12 ip 158.43.192.254
!
```

Connection to Border Routers

```
interface Fddi1/0/0
 description Border FDDI ring for doc-cr1/2 and doc-br1/2
 ip address 158.43.195.1 255.255.255.0
 no ip directed-broadcast
 ip route-cache distributed
 no keepalive
!
```

Core Transit Connection

```
interface FastEthernet1/1/0
 description Transit full-duplex Fast Eth CR1<->CR2
 ip address 158.43.254.65 255.255.255.252
 no ip directed-broadcast
 ip pim sparse-mode
 ip sdr listen
 ip route-cache distributed
 delay 9
 no keepalive
 full-duplex
!
```

STM-1 link to another PoP

```
interface POS4/0/0
 description STM-1 HDLC link to London UK2   M3KL 00367/00  (MFS)
 ip address 158.43.254.25 255.255.255.252
 no ip directed-broadcast
 ip pim sparse-mode
 ip sdr listen
 ip route-cache distributed
 ip ospf cost 1000
 bandwidth 155000
 delay 1000
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
  down-when-looped
  pos framing sdh
  pos flag s1s0 2
 !
```

Connection to Service Routers

```
 interface Fddi5/0/0
  description FDDI ring for Service routers
  ip address 158.43.198.1 255.255.255.0
  no ip directed-broadcast
  ip pim sparse-mode
  ip route-cache distributed
  delay 12
  no keepalive
 !
```

OSPF configuration

```
 router ospf 44
  redistribute connected subnets route-map connected-to-ospf
  redistribute static subnets route-map static-to-ospf
  passive-interface Loopback0
  network 158.43.192.0 0.0.15.255 area 20
  network 158.43.64.0 0.0.15.255 area 20
  network 158.43.254.0 0.0.1.255 area 0
  maximum-paths 3
  default-metric 20
  distance 70
  area 0 range 158.43.254.0 255.255.254.0
  area 20 range 158.43.192.0 255.255.240.0
  area 20 range 158.43.64.0 255.255.240.0
  ospf log-adjacency-changes
 !
```

BGP configuration

```
 router bgp 1849
  no synchronization
  bgp dampening
  network 158.43.0.0 route-map cidr-tag
  redistribute static route-map static-bgp
```

- route reflector configuration for partial and full routes

```
  bgp cluster-id 20
  no bgp client-to-client reflection
  neighbor rr-client peer-group
  neighbor rr-client remote-as 1849
  neighbor rr-client route-reflector-client
  neighbor rr-client update-source Loopback0
  neighbor rr-client send-community
  neighbor rr-client route-map rr-client-out out
  neighbor rr-client-full peer-group
  neighbor rr-client-full remote-as 1849
  neighbor rr-client-full route-reflector-client
  neighbor rr-client-full update-source Loopback0
  neighbor rr-client-full send-community
```

- iBGP config for core network, partial and full routes

```
  neighbor core-ibgp-partial peer-group
  neighbor core-ibgp-partial remote-as 1849
  neighbor core-ibgp-partial update-source Loopback0
  neighbor core-ibgp-partial send-community
  neighbor core-ibgp-partial route-map core-ibgp-partial-out out
  neighbor core-ibgp-full peer-group
  neighbor core-ibgp-full remote-as 1849
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
  neighbor core-ibgp-full update-source Loopback0
  neighbor core-ibgp-full send-community
```

- iBGP peers

```
  neighbor 158.43.131.104 peer-group core-ibgp-full
  neighbor 158.43.131.105 peer-group core-ibgp-full
  neighbor 158.43.162.104 peer-group core-ibgp-full
  neighbor 158.43.162.105 peer-group core-ibgp-full
  neighbor 158.43.179.104 peer-group core-ibgp-full
  neighbor 158.43.179.105 peer-group core-ibgp-full
  neighbor 158.43.206.3 peer-group rr-client
  neighbor 158.43.206.4 peer-group rr-client
  neighbor 158.43.206.5 peer-group rr-client
  neighbor 158.43.206.6 peer-group rr-client
  neighbor 158.43.206.7 peer-group rr-client
  neighbor 158.43.206.8 peer-group rr-client
  neighbor 158.43.206.9 peer-group rr-client
  neighbor 158.43.206.10 peer-group rr-client
  neighbor 158.43.206.11 peer-group rr-client
  neighbor 158.43.206.12 peer-group rr-client
  neighbor 158.43.206.13 peer-group rr-client
  neighbor 158.43.206.14 peer-group rr-client
  neighbor 158.43.206.15 peer-group rr-client
  neighbor 158.43.206.64 peer-group rr-client
  neighbor 158.43.206.65 peer-group rr-client
  neighbor 158.43.206.66 peer-group rr-client
  neighbor 158.43.206.84 peer-group rr-client
  neighbor 158.43.206.96 peer-group core-ibgp-partial
  neighbor 158.43.206.97 peer-group core-ibgp-partial
  neighbor 158.43.206.103 peer-group core-ibgp-partial
  neighbor 158.43.206.105 peer-group core-ibgp-full
  neighbor 158.43.211.104 peer-group core-ibgp-full
  neighbor 158.43.211.105 peer-group core-ibgp-full
  neighbor 158.43.219.104 peer-group core-ibgp-full
  neighbor 158.43.219.105 peer-group core-ibgp-full
  neighbor 158.43.227.104 peer-group core-ibgp-full
  neighbor 158.43.227.105 peer-group core-ibgp-full
  neighbor 158.43.234.96 peer-group core-ibgp-full
  neighbor 158.43.234.104 peer-group core-ibgp-full
  neighbor 158.43.234.105 peer-group core-ibgp-full
  neighbor 158.43.239.104 peer-group core-ibgp-full
  neighbor 158.43.239.105 peer-group core-ibgp-full
  neighbor 158.43.251.28 peer-group core-ibgp-full
  neighbor 158.43.251.29 peer-group core-ibgp-full
```

- and the rest…

```
  maximum-paths 2
  distance bgp 180 200 200
  no auto-summary
 !
```

Static routes for network blocks

```
 ip route 193.128.0.0 255.252.0.0 Null0
 ip route 193.132.0.0 255.254.0.0 Null0
 ip route 194.128.0.0 255.252.0.0 Null0
 ip route 194.200.0.0 255.252.0.0 Null0
 ip route 194.216.0.0 255.255.0.0 Null0
 ip route 195.217.0.0 255.255.0.0 Null0
 !
```

Static routes to legacy Access Servers and other Server systems

```
 ip route 158.43.70.0 255.255.255.192 158.43.192.250
```

```
ip route 158.43.82.64 255.255.255.192 158.43.192.250
ip route 158.43.86.0 255.255.255.192 158.43.192.250
ip route 158.43.86.64 255.255.255.192 158.43.192.250
ip route 158.43.86.128 255.255.255.192 158.43.192.250
ip route 158.43.86.192 255.255.255.192 158.43.192.251
ip route 158.43.240.4 255.255.255.255 158.43.192.1
ip route 159.100.0.0 255.255.0.0 158.43.192.251
ip route 193.36.1.0 255.255.255.0 158.43.192.251
…<snip>
!
route-map cidr-tag permit 10
 match ip address 143
 set community 1849:5005
!
route-map rr-client-out permit 10
 match community 1 10 12 18 22
!
…<snip>
!
end
```

Gateway Router

This section looks at a typical gateway router configuration – this router is connected to the core router discussed above by FastEthernet. Configuration which has already been discussed previously has been omitted for the sake of brevity.

```
!
hostname doc-gw4
!
<snip>
!
```

Channelised E1 controller configuration

- 64K timeslots

```
controller E1 5/0
 framing NO-CRC4
 channel-group 0 timeslots 1
 channel-group 1 timeslots 2
 <snip>
 channel-group 28 timeslots 29
 channel-group 29 timeslots 30
 description BT - GXUK xxxxxx
!
```

- nx64K timeslots

```
controller E1 5/1
 framing NO-CRC4
 channel-group 0 timeslots 1-4
 channel-group 4 timeslots 5-6
 channel-group 6 timeslots 7-8
 channel-group 8 timeslots 9-10
 channel-group 10 timeslots 11-12
 channel-group 12 timeslots 13-20
 channel-group 20 timeslots 21-24
 channel-group 24 timeslots 25-28
 channel-group 28 timeslots 29-30
 description BT - GXUK yyyyyy
!
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

- 64K timeslots

```
controller E1 6/0
  <snip>
!
controller E1 6/1
  <snip>
!
```

Loopbacks never go away!

```
interface Loopback0
 ip address 158.43.206.4 255.255.255.255
!
```

Connection to Backbone GW router LAN

```
interface FastEthernet0/0/0
 description Docklands PoP Core Fast Ethernet
 ip address 158.43.200.4 255.255.255.0
 no ip redirects
 no ip directed-broadcast
 ip route-cache flow
 no ip route-cache optimum
 ip route-cache distributed
 delay 15
 full-duplex
!
```

FSIP connected customers – note description, circuit ID and cable number!

```
interface Serial1/0
 description HDLC link to ebscpubld1-1  NXUK abcdef  DA0
 ip unnumbered Loopback0
 bandwidth 128
!
interface Serial1/1
 description link to Pinnacle Internet Services Ltd, Station Way  NXUK ghijkl  DA1
 ip unnumbered Loopback0
 bandwidth 128
!
interface Serial1/2
 description HDLC primary link to IBM PC User Group, Harrow  NXUK mnopqrs  DA2
 ip address 158.43.65.161 255.255.255.252
 bandwidth 512
!
<snip>
!
```

MIP2 connected 64k customer links – note details

```
interface Serial5/0:29
 description HDLC link to Copyright Licensing Agency,   KXUK aaaaaa
 ip unnumbered Loopback0
 bandwidth 64
 transmit-buffers backing-store
!
interface Serial5/0:28
 description HDLC link to highfield1-1  KXUK bbbbbb
 ip unnumbered Loopback0
 bandwidth 64
 transmit-buffers backing-store
!
<snip>
!
```

MIP2 connected nx64k customer links – note details

```
<snip>
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
interface Serial5/1:20
 description 256k HDLC link to Text 100 Limited [text100]  NXUK xxx123
 ip unnumbered Loopback0
 no ip mroute-cache
 bandwidth 256
 transmit-buffers backing-store
!
interface Serial5/1:12
 description 512k HDLC link to Theodore Goddard [theogodd1-1]  NXUK zzz456
 ip unnumbered Loopback0
 no ip mroute-cache
 bandwidth 512
 transmit-buffers backing-store
!
<snip>
!
```

Historic link – X25 enacapsulation!

```
interface Serial6/1:29
 description X.25 link to Fisons Surface Systems, East Grinstead  KXUK abc987
 ip address 158.43.66.193 255.255.255.252
 no ip mroute-cache
 encapsulation x25 dce
 bandwidth 64
 x25 address 000008400060
 x25 htc 4
 x25 win 7
 x25 wout 7
 x25 ips 1024
 x25 ops 1024
 x25 map ip 158.43.66.194 000008400061 packetsize 1024 1024 windowsize 7 7 nvc 4
 transmit-buffers backing-store
!
<snip>
!
```

OSPF configuration

```
router ospf 44
 redistribute connected subnets route-map connected-to-ospf
 redistribute static subnets route-map static-to-ospf
 passive-interface Serial1/0
 <snip>
 passive-interface Loopback0
 network 158.43.192.0 0.0.15.255 area 20
 network 158.43.64.0 0.0.15.255 area 20
 maximum-paths 6
 default-metric 20
 distance 70
!
```

BGP configuration

```
router bgp 1849
 no synchronization
 bgp dampening
 network 192.68.174.0 route-map static-bgp
 network 192.86.127.0 route-map static-bgp
 redistribute static route-map static-bgp
 neighbor reflector peer-group
 neighbor reflector remote-as 1849
 neighbor reflector update-source Loopback0
 neighbor reflector send-community
 neighbor client-peer peer-group
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
 neighbor client-peer remote-as 1849
 neighbor client-peer update-source Loopback0
 neighbor client-peer send-community
```

- customer multihomed on AS1849 – see interface serial 1/2 above – note default originate

```
 neighbor 158.43.65.162 remote-as 2830
 neighbor 158.43.65.162 default-originate
 neighbor 158.43.65.162 distribute-list 101 in
 neighbor 158.43.65.162 distribute-list 168 out
 neighbor 158.43.65.162 route-map uk-cust-in in
 neighbor 158.43.65.162 route-map defaultroute-out out
 neighbor 158.43.65.162 filter-list 40 in
```

- iBGP peering with rest of RR cluster routers

```
 neighbor 158.43.206.3 peer-group client-peer
 neighbor 158.43.206.5 peer-group client-peer
 neighbor 158.43.206.6 peer-group client-peer
 neighbor 158.43.206.7 peer-group client-peer
 neighbor 158.43.206.8 peer-group client-peer
 neighbor 158.43.206.9 peer-group client-peer
 neighbor 158.43.206.10 peer-group client-peer
 neighbor 158.43.206.11 peer-group client-peer
 neighbor 158.43.206.12 peer-group client-peer
 neighbor 158.43.206.13 peer-group client-peer
 neighbor 158.43.206.14 peer-group client-peer
 neighbor 158.43.206.15 peer-group client-peer
 neighbor 158.43.206.64 peer-group client-peer
 neighbor 158.43.206.65 peer-group client-peer
 neighbor 158.43.206.66 peer-group client-peer
 neighbor 158.43.206.84 peer-group client-peer
 neighbor 158.43.206.104 peer-group reflector
 neighbor 158.43.206.105 peer-group reflector
```

- miscellaneous BGP

```
 distance bgp 180 200 200
 no auto-summary
!
<snip>
!
```

Static routes for IP unnumbered links and others – note IP unnumbered!

```
<snip>
ip route 193.128.85.32 255.255.255.224 Serial5/0:28
ip route 193.130.56.0 255.255.252.0 158.43.66.194
ip route 193.130.60.0 255.255.255.0 158.43.66.194
ip route 193.130.63.48 255.255.255.240 Serial5/1:12
ip route 193.132.233.192 255.255.255.224 Serial1/0
ip route 193.133.228.16 255.255.255.240 Serial5/0:28
ip route 194.128.198.0 255.255.254.0 Serial1/1
ip route 194.128.212.0 255.255.255.0 Serial5/0:29
ip route 194.201.30.0 255.255.255.0 Serial5/1:20
<snip>
```

AS Path access lists

```
ip as-path access-list 2 permit ^(2830_)+$
ip as-path access-list 2 deny .*
<snip>
ip as-path access-list 40 permit ^(_[0-9]+)$
ip as-path access-list 40 deny .*
<snip>
```

Standard Access Lists

```
access-list 101 permit ip host 193.128.16.0 host 255.255.255.0
access-list 101 permit ip host 193.128.17.0 host 255.255.255.0
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

```
access-list 101 permit ip host 192.68.174.0 host 255.255.255.0
access-list 101 permit ip host 192.86.127.0 host 255.255.255.0
access-list 101 permit ip host 194.202.72.0 host 255.255.255.0
access-list 101 deny   ip any any
<snip>
```

route map applied to all BGP peering customers (big + detailed!)

```
route-map uk-cust-in permit 10
 match community 6
!
route-map uk-cust-in permit 20
 match ip address 144
 set community 1849:5666 local-AS additive
!
route-map uk-cust-in permit 30
 match community 7
 set local-preference 70
 set community 1849:5666 additive
!
route-map uk-cust-in permit 40
 match community 8
 set local-preference 80
 set community 1849:5666 additive
!
route-map uk-cust-in permit 50
 match community 9
 set local-preference 90
 set community 1849:5666 additive
!
route-map uk-cust-in permit 60
 match community 11
 set local-preference 110
 set community 1849:5666 additive
!
route-map uk-cust-in permit 70
 match community 13
 set local-preference 130
 set community 1849:5666 additive
!
route-map uk-cust-in permit 80
 match community 23 24 25 26 27 28
 set community 1849:5666 additive
!
route-map uk-cust-in permit 90
 set community 1849:5666 additive
!
```

send default route only to customer

```
route-map defaultroute-out deny 20
 match community 21
!
```

rest similar to previously…

```
<snip>
end
```


Cisco Systems, Inc.
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100