Studio del traffico con Netflow

Massimo Carboni - Direzione GARR-B
Massimo.Carboni@garr.it
III WorkShop GARR-B
Firenze
24-25 Gennaio 2001

Indice

- Network Design e Sicurezza informatica
- Configurazione Router
  - Analisi dal Router
- Soluzione software: Cflowd
- Strumenti di analisi: Arts Tools
- La configurazione in GARR-B

Configurazione del Router

```
RT_NAPOLI# configure terminal
RT_NAPOLI(config)#ip flow-export destination 193.206.158.20 7777
RT_NAPOLI(config)#ip flow-export source Loopback0
RT_NAPOLI(config)#ip flow-export version 5
RT_NAPOLI(config)#int atm 1/0/0
RT_NAPOLI(config-if)#ip route-cache flow
RT_NAPOLI(config-if)#exit
RT_NAPOLI#exit
RT_NAPOLI>sh ip flow export
Flow export is enabled
  Exporting flows to 193.206.158.20 (7777)
  Exporting using source IP address Loopback0
 Version 5 flow records, origin-as
  295168459 flows exported in 9838956 udp datagrams
 ...
```

Sh ip cache flow (1/2)

```
RT_NAPOLI>sh ip cache flow
IP packet size distribution (60127M total packets):
1-32   64   96  128  160  192  224  256  288  320  352  384  416  448  480
.002 .324 .032 .010 .010 .007 .006 .006 .005 .005 .004 .006 .004 .003 .004

 512  544  576 1024 1536 2048 2560 3072 3584 4096 4608
.003 .003 .088 .071 .395 .000 .000 .000 .000 .000 .000

IP Flow Switching Cache, 22013888 bytes
  14 active, 311166 inactive, 3073403967 added
  2915254561 ager polls, 0 flow alloc failures
  Active flows timeout in 10 minutes
  last clearing of statistics 22w2d
```

Sh ip cache flow (2/2)

Protocol	Total Flows	Flows /Sec	Packets /Flow	Bytes /Pkt	Packets /Sec	Active(Sec) /Flow	Idle(Sec) /Flow
TCP-Telnet	11540345	2.6	9	113	26.6	5.1	18.5
TCP-FTP	51652795	12.0	5	129	70.2	5.9	14.2
TCP-FTPD	18033723	4.1	223	1009	939.6	31.1	13.2
TCP-WWW	1806622735	420.6	15	834	6339.7	5.0	20.4
TCP-SMTP	87637272	20.4	9	335	201.5	6.8	15.7
TCP-X	243337	0.0	180	403	10.2	33.1	16.9
TCP-BGP	372703	0.0	2	62	0.2	9.4	15.6
TCP-NNTP	6854529	1.5	173	1079	276.9	21.5	15.1
TCP-Frag	55991	0.0	28	165	0.3	15.3	28.0
TCP-other	1004687740	233.9	21	641	5062.1	6.5	16.7
UDP-DNS	398716016	92.8	2	104	262.7	4.8	23.6
UDP-NTP	16239673	3.7	1	76	4.3	1.3	21.2
UDP-TFTP	9471	0.0	7	157	0.0	2.0	21.3
UDP-Frag	2994674	0.6	22	845	15.7	3.2	14.4
UDP-other	337405126	78.5	6	402	486.8	3.2	25.2
ICMP	234215176	54.5	5	168	296.3	6.5	23.2
IGMP	8797	0.0	124	1286	0.2	2.8	16.7
IPINIP	231002	0.0	6	173	0.3	16.8	19.7
GRE	91194	0.0	82	232	1.7	220.6	10.7
IP-other	4607365	1.0	2	85	3.2	8.7	16.6
Total:	3982219664	927.1	15	725	13999.3	5.5	20.1

Struttura dei dati (1/2)

byte 3	byte 2	byte 1	byte 0

Version 5 Flow Header

version		count
sysUptime		
unix seconds		
unix nanoseconds		
flow sequence number		
engine type	engine ID	reserved

Struttura dei dati (2/2)

byte 3	byte 2	byte 1	byte 0

Version 5 Flow Entry

source IP address	
destination IP address	
next hop IP address	
input interface index	output interface index
packets	
bytes	
start time of flow	
end time of flow	
source port	destination port
pad / TCP flags	IP protocol / TOS
source AS	destination AS
src netmask length / dst netmask length	padding

Arts + Cflowd

http://www.caida.org/

- Gira su Linux (RH-6.2)
  - richiede il compilatore GCC-2.95
- Arts library (v-0.9.b6)
  - DOC: http://www.caida.org/tools/utilities/arts/
  - SOFT: ftp://ftp.caida.org/pub/arts++/arts++-0-9-b6.tar.gz
- Cflowd
  - DOC: http://www.caida.org/tools/measurement/cflowd/
  - SOFT: ftp://ftp.caida.org/pub/cflowd/cflowd-2-1-b1.tar.gz

Architettura Software: Cflowd

- Livello di aggregazione
  - *cflowdmux* + *cflowd*
  - Flusso UDP dai Router
- Salvataggio dei dati
  - *cflowd* Þ *cfdcollect*
  - Flusso TCP dai differenti processi
- Analisi dei dati:
  - arts tools

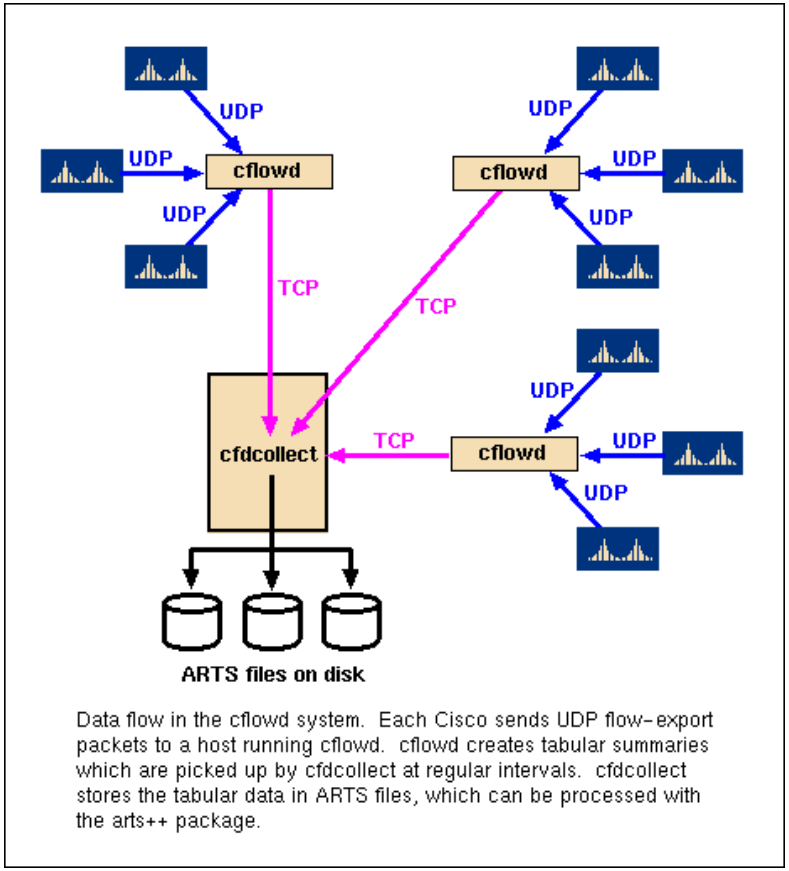

Data flow in the cflowd system. Each Cisco sends UDP flow-export packets to a host running cflowd. cflowd creates tabular summaries which are picked up by cfdcollect at regular intervals. cfdcollect stores the tabular data in ARTS files, which can be processed with the arts++ package.

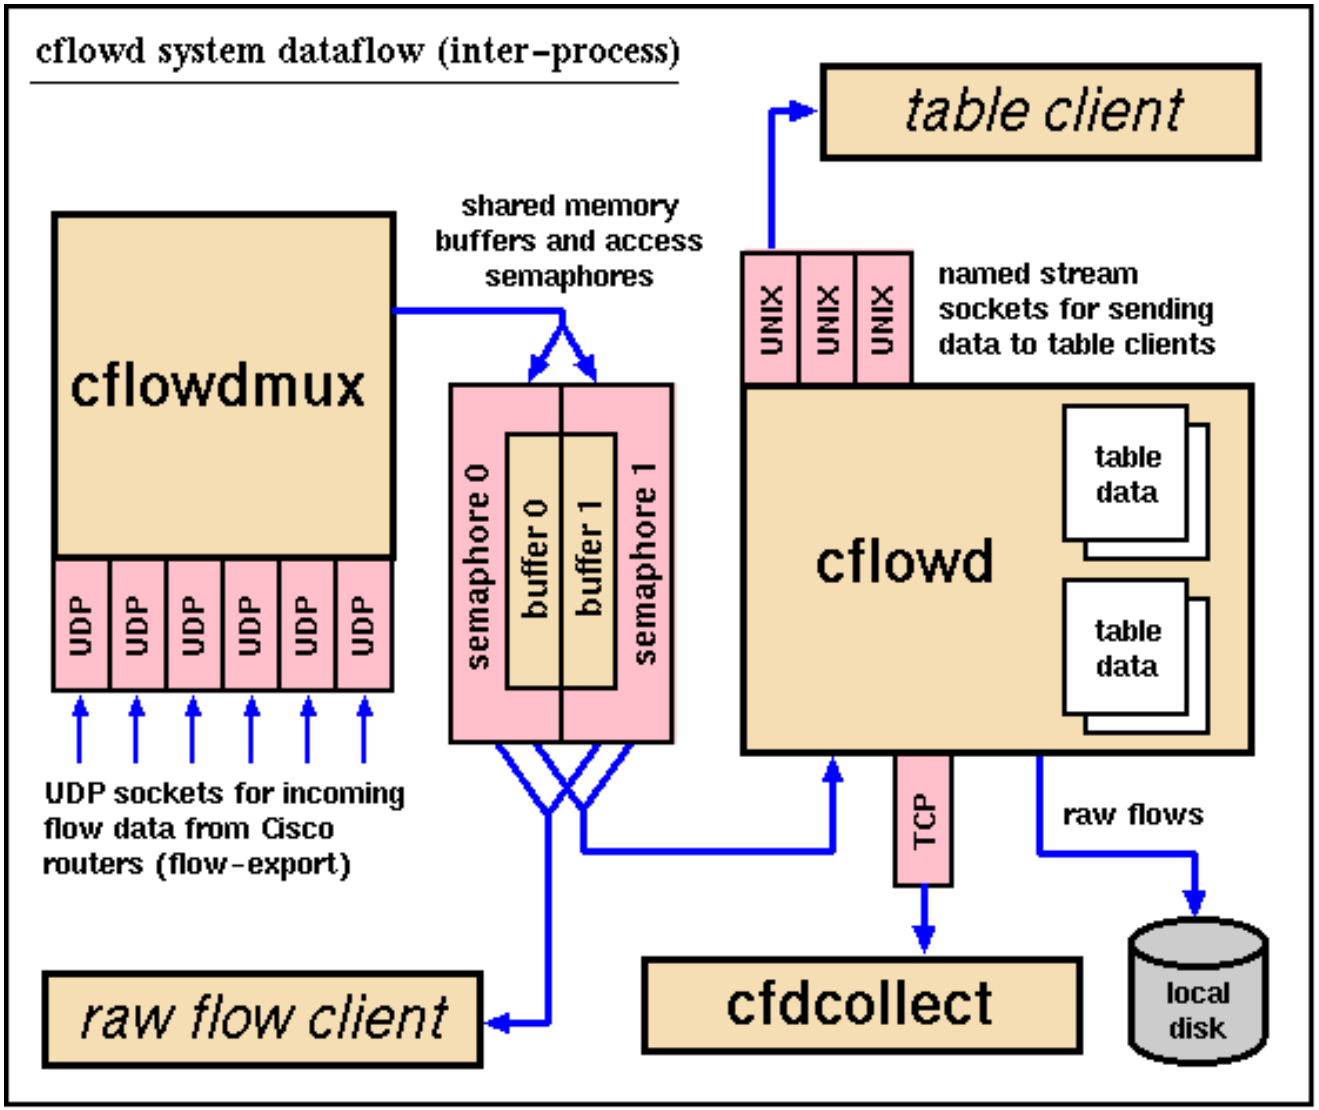
cflowd system dataflow (inter-process)
table client
shared memory
buffers and access
semaphores
cflowdmux
UNIX
UNIX
UNIX
named stream
sockets for sending
data to table clients
UDP
UDP
UDP
UDP
UDP
UDP
semaphore 0
buffer 0
buffer 1
semaphore 1
cflowd
table
data
table
data
UDP sockets for incoming
flow data from Cisco
routers (flow-export)
TCP
raw flows
raw flow client
cfdcollect
local
disk

Configurazione Cflowd (1/2)

```
#File: /usr/local/arts/etc/cflowd.conf
#-----------------------------------------------------------
OPTIONS {
  LOGFACILITY:          local6
  TCPCOLLECTPORT:       2056
  PKTBUFSIZE:           2097152
  TABLESOCKFILE:        /usr/local/arts/etc/cflowdtable.socket
  FLOWDIR:              /usr/local/arts/data/cflowd/flows
  FLOWFILELEN:          2097152
  NUMFLOWFILES:         10
  MINLOGMISSED:         300
}
#-----------------------------------------------------------
COLLECTOR {
  HOST:         193.206.158.31  # IP address of central collector
  ADDRESSES:    { 193.206.158.31 }
  AUTH:         none
}
```

Configurazione Cflowd (2/2)

```
#File: /usr/local/arts/etc/cflowd.conf
#-------------------------------------------------------
CISCOEXPORTER {
  HOST: 193.206.129.252 # IP address of central collector
 ADDRESSES: {
   193.206.135.4,   # Ethernet0
    212.1.200.26,   # POS3/0
   193.206.134.1,   # ATM4/0.101
  193.206.134.17,   # ATM5/0.100
   193.206.134.9,   # ATM6/0.103
 193.206.134.210}   # FEth8/0
  CFDATAPORT: 8001
    SNMPCOMM: 'public'
     LOCALAS: 137       #  Local AS of Cisco sending data.
     COLLECT: { protocol, portmatrix,ifmatrix, nexthop,
                netmatrix, asmatrix, tos, flows }
}
```

Configurazione Cfdcollect

```
#File: /usr/local/arts/etc/cfdcollect.conf
#-------------------------------------------------------------
system {
  logFacility:          local6          #  Syslog to local6 facility.
  dataDirectory:        /usr/local/arts/data/cflowd
  filePrefix:           arts
  pidFile:              /usr/local/arts/etc/cfdcollect.pid
}
#-------------------------------------------------------------
cflowd {
  host:                 193.206.135.215
  tcpCollectPort:       2056
  minPollInterval:      600
}
```

Configurazione Syslog

```
# Modificare il file: /etc/syslog.conf
                .
                .
                .
local6.*                        /var/log/cflowd.log
                .
                .
                .

# File: cflowd_start.sh
# Esecuzione come utente non privilegiato.
 /usr/local/arts/sbin/cflowdmux
 /usr/local/arts/sbin/cflowd
 /usr/local/arts/sbin/cfdcollect \
        /usr/local/arts/etc/cflowdcollect.conf
#
```

Arts Tools

```
set path = ( /usr/local/arts/bin $path)

setenv ARTS_ROOT        /usr/local/arts
setenv ARTS_MILANO_RT   /usr/local/arts/data/cflowd/193.206.129.254
setenv ARTS_MILANO2_RT  /usr/local/arts/data/cflowd/193.206.129.252
setenv ARTS_BOLOGNA_RT  /usr/local/arts/data/cflowd/193.206.128.254
setenv ARTS_ROMA_RT     /usr/local/arts/data/cflowd/193.206.131.254
setenv ARTS_NAPOLI_RT   /usr/local/arts/data/cflowd/193.206.130.254
setenv ARTS_MIX         /usr/local/arts/data/cflowd/193.206.134.250
setenv ARTS_RIX         /usr/local/arts/data/cflowd/193.206.134.226

setenv MANPATH          /usr/local/arts/man:$MANPATH
```

artsprotos

```
# artsprotos -i 30 $ARTS_MILANO2_RT/arts.20010122
  router:  193.206.129.252
  ifIndex: 30
  period:  01/22/2001 15:00:25 - 01/22/2001 15:10:28 CET
  Protocol            Pkts       Pkts/sec           Bytes       Bits/sec
  --------  -------------  -------------  -------------  -------------
       tcp       14401938          23883    13061644575      173288816
       udp         469243            778      161996274        2149204
      icmp         177971            295       12950755         171817
        93            338              0          30472            404
       110             17              0          17748            235
   ipencap             70              0           6181             82
ipv6-crypt             36              0           5512             73
      ipv6             19              0           1596             21
```

artsnets

```
# artsnets -b '01/22/2001 15:00:25' -i 30 $ARTS_MILANO2_RT/arts.20010122
  router:  193.206.129.252
 ifIndex: 30
  period:  01/22/2001 15:00:25 - 01/22/2001 15:10:28 CET
         Src Network          Dst Network            Pkts           Bytes
  -----------------  ------------------  -------------  -------------
  206.204.210.198/32     192.107.81.0/24         269081      381833893
      131.225.0.0/16    193.205.157.0/24         141251      204633304
    206.204.7.100/32      131.175.0.0/16          96897      137521507
    206.204.7.116/32      159.149.0.0/16          91451      137141200
     207.189.64.0/19     193.204.64.0/20          81000      121474716
       134.79.0.0/16      140.105.0.0/16          77803      115310011
      165.230.0.0/16      131.114.0.0/17          89490      100308743
     130.227.83.9/32    193.204.176.0/20          66045       98412218
     209.114.71.9/32       160.78.0.0/16          65401       96592892
    208.39.14.158/32      157.138.0.0/16          63654       94743431
      157.182.0.0/16      137.204.0.0/16          63263       93546501
```

artsports

```
# artsportmagg -s1-65535 -i 30 /tmp/ports.mi2ny.20010122 \
                        $ARTS_MILANO2_RT/arts.20010122
# artsports /tmp/ports.mi2ny.20010122
router:  193.206.129.252
ifIndex: 30
period:  01/22/2001 00:50:28 - 01/23/2001 00:50:30 CET
selected ports: 1-65535
```

Port	InPkts	InBytes	OutPkts	OutBytes
www	18419261	1262716424	391344267	365134378963
ftp-data	19647911	11512669526	28531273	34637352684
6699	10263643	7690716540	5941295	6160192545
nntp	7724108	9965403565	3666142	3767448891
ssh	3750981	4924872408	15653277	3890106916
6346	14793863	5942318651	1404418	601934441
5501	10899	437849	5100738	5783006256
smtp	8320369	4683453429	5462881	344415891
6688	1308206	712908008	3083537	3247187880

FlowDump

```
FLOW
  index:           0xc7ffff
  router:          193.206.131.254
  src IP:          193.204.44.80
  dst IP:          212.171.4.180
  input ifIndex:   101
  output ifIndex:  116
  src port:        80
  dst port:        1245
  pkts:            9
  bytes:           4170
  IP nexthop:      193.206.134.226
  start time:      Mon Jan 22 18:16:43 2001
  end time:        Mon Jan 22 18:16:44 2001
  protocol:        6
  tos:             0
  src AS:          137
  dst AS:          137
  src masklen:     20
  dst masklen:     19
  TCP flags:       0x1b
  engine type:     1
  engine id:       4
```

Lettura dei dati

```
#!/usr/bin/perl
$flodump='/usr/local/arts/bin/flowdump';
$datadir='/usr/local/arts/data/cflowd/flows/193.206.134.252.flows.*';
open (FLOW,"$flowdump $datadir|");
while (<FLOW>){
    chomp;
    if ( /FLOW/ ) {$Bytes = 0;$SPort = 0;$SHost = '';}
    $Bytes = $1 if ( /bytes:\s*(\d*)/ );
    $SHost = $1 if ( /src IP:          (\d+.\d+.\d+.\d+)/ );
    $SPort = $1 if ( /src port:\s+(\d+)$/ ) ;
    $TotBytes += $Bytes;
    $Bytes {$SHost} += $Bytes;
}
foreach $HOST ( sort keys %Bytes ) {
    printf ("HOST %15s :: Bytes :: %12.1f :: Perc %5.2f%%\n",
       $HOST,$Bytes{$HOST},$Bytes{$HOST}/$TotBytes*100)
       if ($Bytes{$HOST}/$TotBytes*100 > 1);
}
print "Total Bytes: $TotBytes\n";
```

Implementazione in GARR-B

PoP di Frascati

Aggregazione flussi UDP
- 5 router di trasporto
- 2 router di peering

Pentium II 400 Mhz
512 MB/RAM
5Mbps di flusso

Direzione GARR-B

Salvataggio dei dati aggregati per analisi off-line

Dual Processor PIII 600 Mhz
1GB RAM 60 GB RAID disk
0.5Mbps di flusso

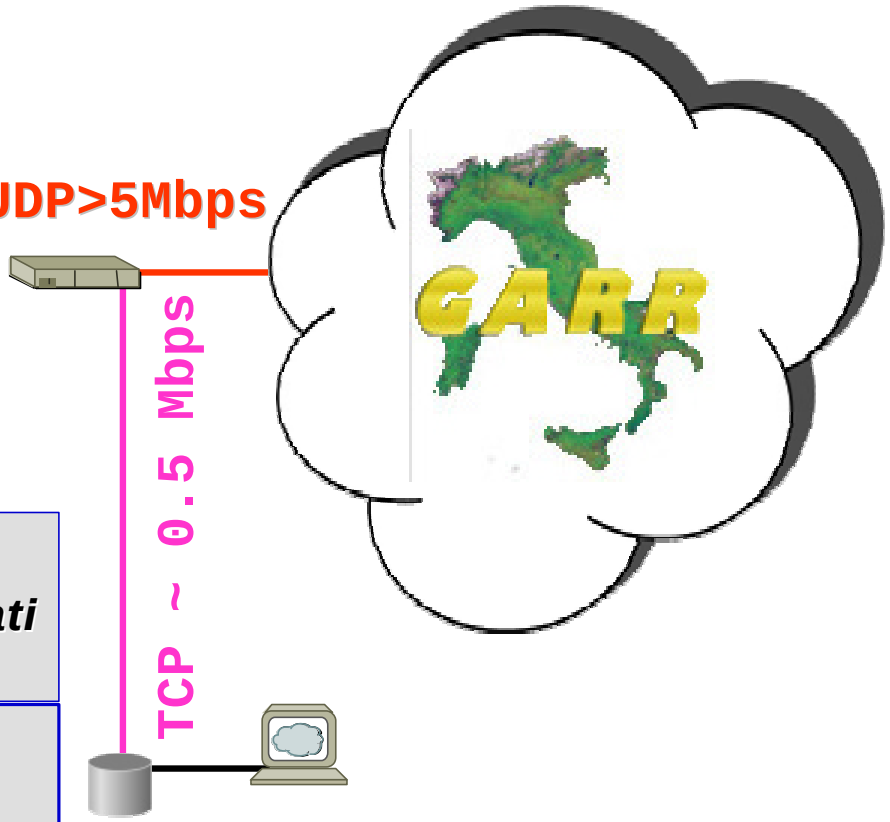

Carico sui server

```
 2:50pm  up 6 days,  4:58,  3 users,  load average: 1.41, 0.96, 0.79
42 processes: 36 sleeping, 5 running, 0 zombie, 1 stopped
CPU states: 43.5% user, 24.6% system,  0.0% nice, 31.7% idle
Mem:   517224K av,  515328K used,    1896K free,   28084K shrd,   27772K buff
Swap:  789312K av,    7160K used,  782152K free                   33288K cached

  PID USER     PRI  NI  SIZE  RSS SHARE STAT  LIB %CPU %MEM   TIME COMMAND
 5687 carboni   18   0  225M 223M 15724 R       0 42.6 44.1  10:28 cflowd
 5768 carboni    6   0  210M 210M   420 R       0 20.5 41.7   0:08 cflowd
 5676 carboni    1   0  8676 8660  8564 R       0  2.3  1.6   0:59 cflowdmux

 2:52pm  up 42 days,  2:27,  5 users,  load average: 0.11, 0.06, 0.31
45 processes: 44 sleeping, 1 running, 0 zombie, 0 stopped
CPU states:  0.0% user,  0.6% system,  0.0% nice, 99.3% idle
Mem:   517180K av,  483724K used,   33456K free,   10600K shrd,  182932K buff
Swap:  526296K av,    3612K used,  522684K free                  207532K cached

  PID USER     PRI  NI  SIZE  RSS SHARE STAT  LIB %CPU %MEM   TIME COMMAND
 6657 carboni   17   0 69240  67M   960 S       0  0.7 13.3   0:10 cfdcollect
```

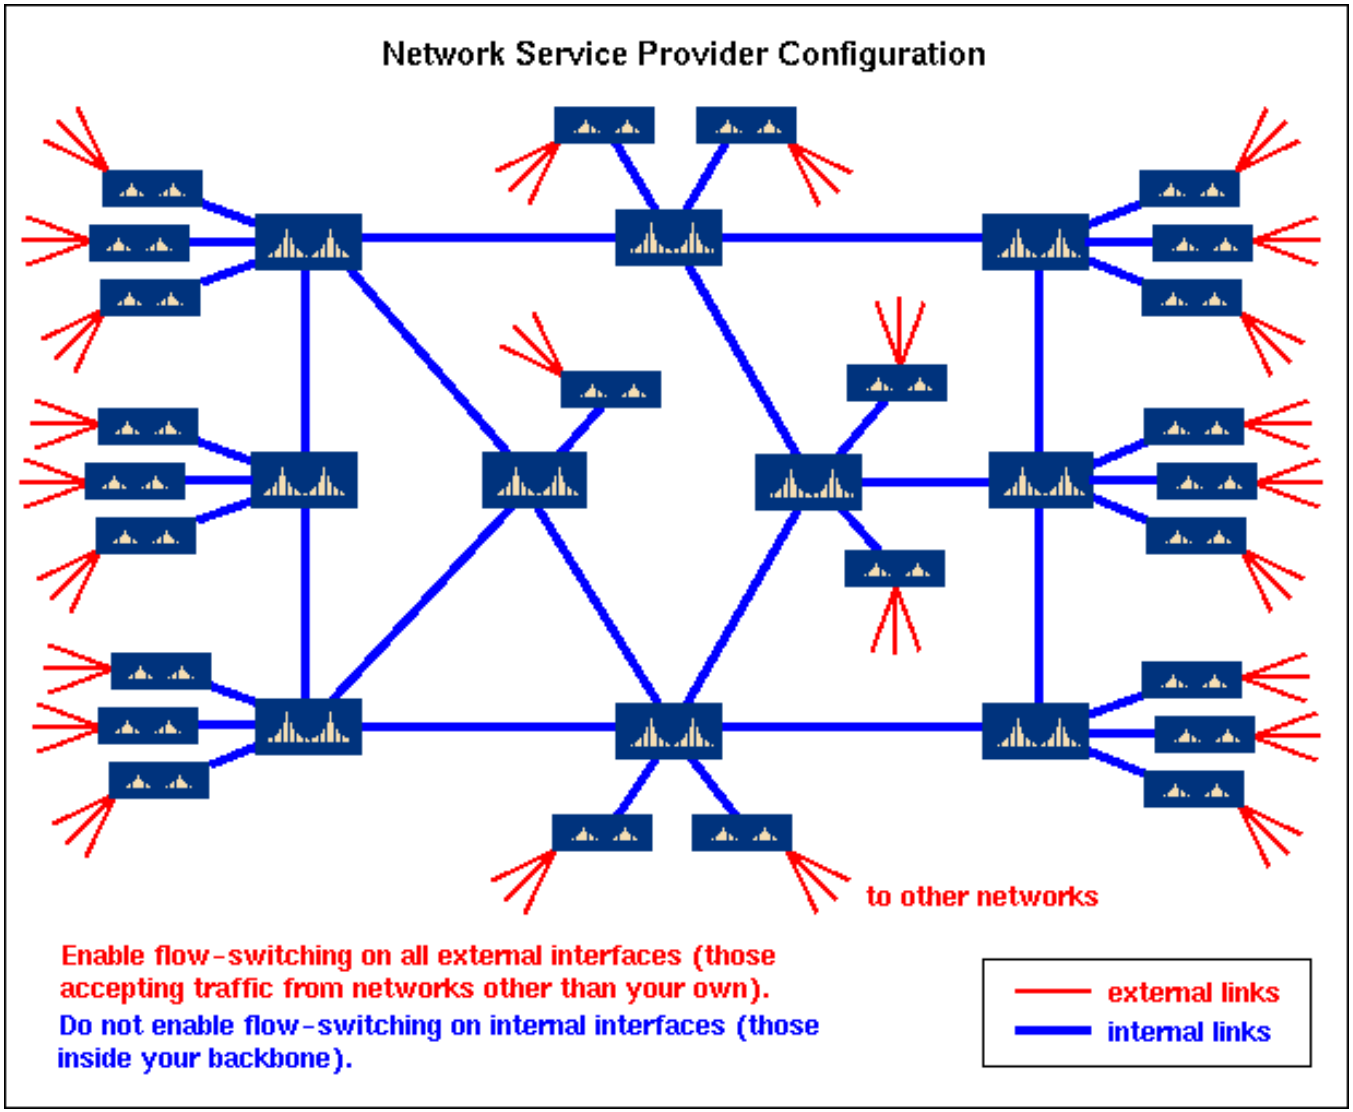
Network Service Provider Configuration
to other networks
Enable flow-switching on all external interfaces (those accepting traffic from networks other than your own).
Do not enable flow-switching on internal interfaces (those inside your backbone).
external links
internal links

Alcuni puntatori

- **http://www.caida.org/**
- **ftp://ftp-eng.cisco.com/ftp/drowell/flow_agg.pdf**
- **http://www.caida.org/tools/utilities/arts/**
- **ftp://ftp.caida.org/pub/arts++/arts++-0-9-b6.tar.gz**
- **http://www.caida.org/tools/measurement/cflowd/**
- **ftp://ftp.caida.org/pub/cflowd/cflowd-2-1-b1.tar.gz**

That's all folks