IP in the mobile core

NANOG 47

October 20, 2009

Jan Chrillesen chrille@tdc.net - TDC A/S

Background

- TDC – you know as AS 3292 but we also do voice!

- TDC mobile seperate company until recently
- Transmission and IP through wholesale from TDC
- 3 years ago all IP was handed over to TDC operations
- 1 year ago all mobile network operations outsourced to Ericsson
- Part of oursourcing contract – to build new combined 2G/3G core
- Part of outsourcing contract – core transmission to use IP and move most/all other interfaces to IP

How is this mobile thing different from IP?

- Everything regarding mobile networks seems to be a very closed world
- Lot of reliance on vendors
- Not really any public mailinglists
- Conferences seems to be mostly focused on the business side
- Most people still thinks of circuits
- Userplane and control plane traffic is different interfaces
- A phone is always referred to as a terminal!

Introduction to GSM and UMTS networks

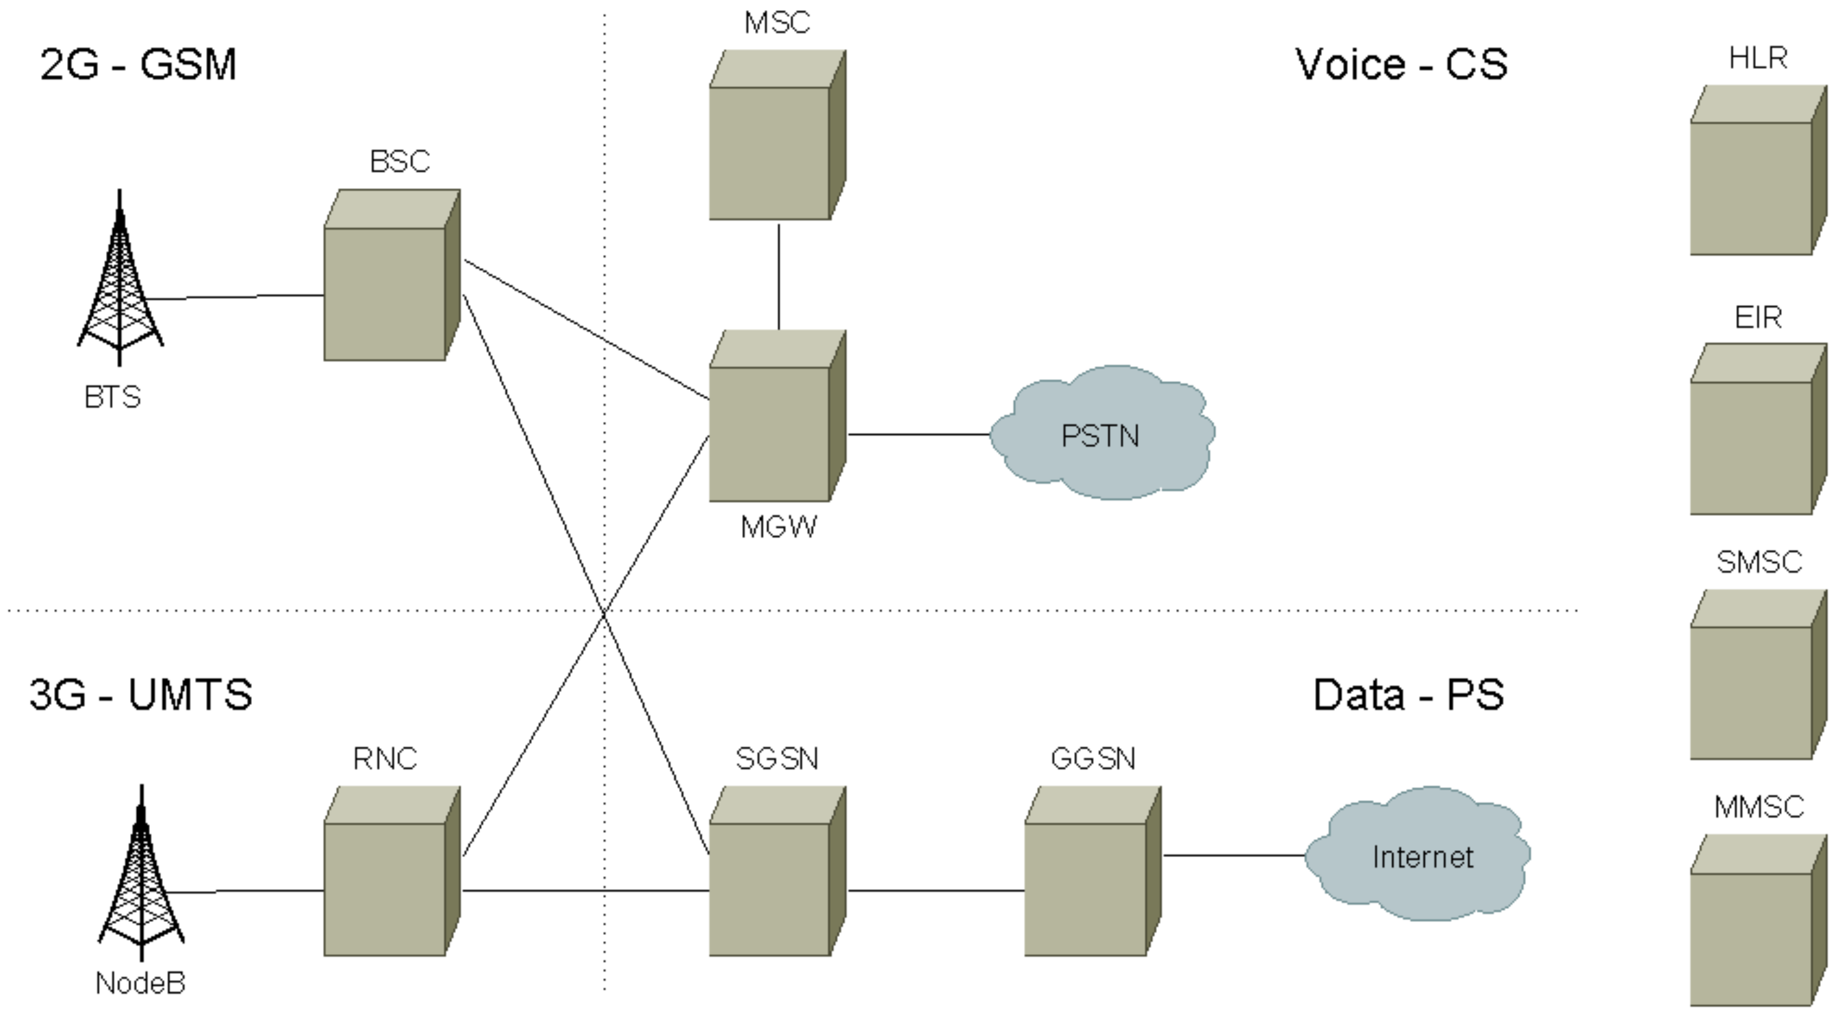

Scale of TDC network and legacy transmission

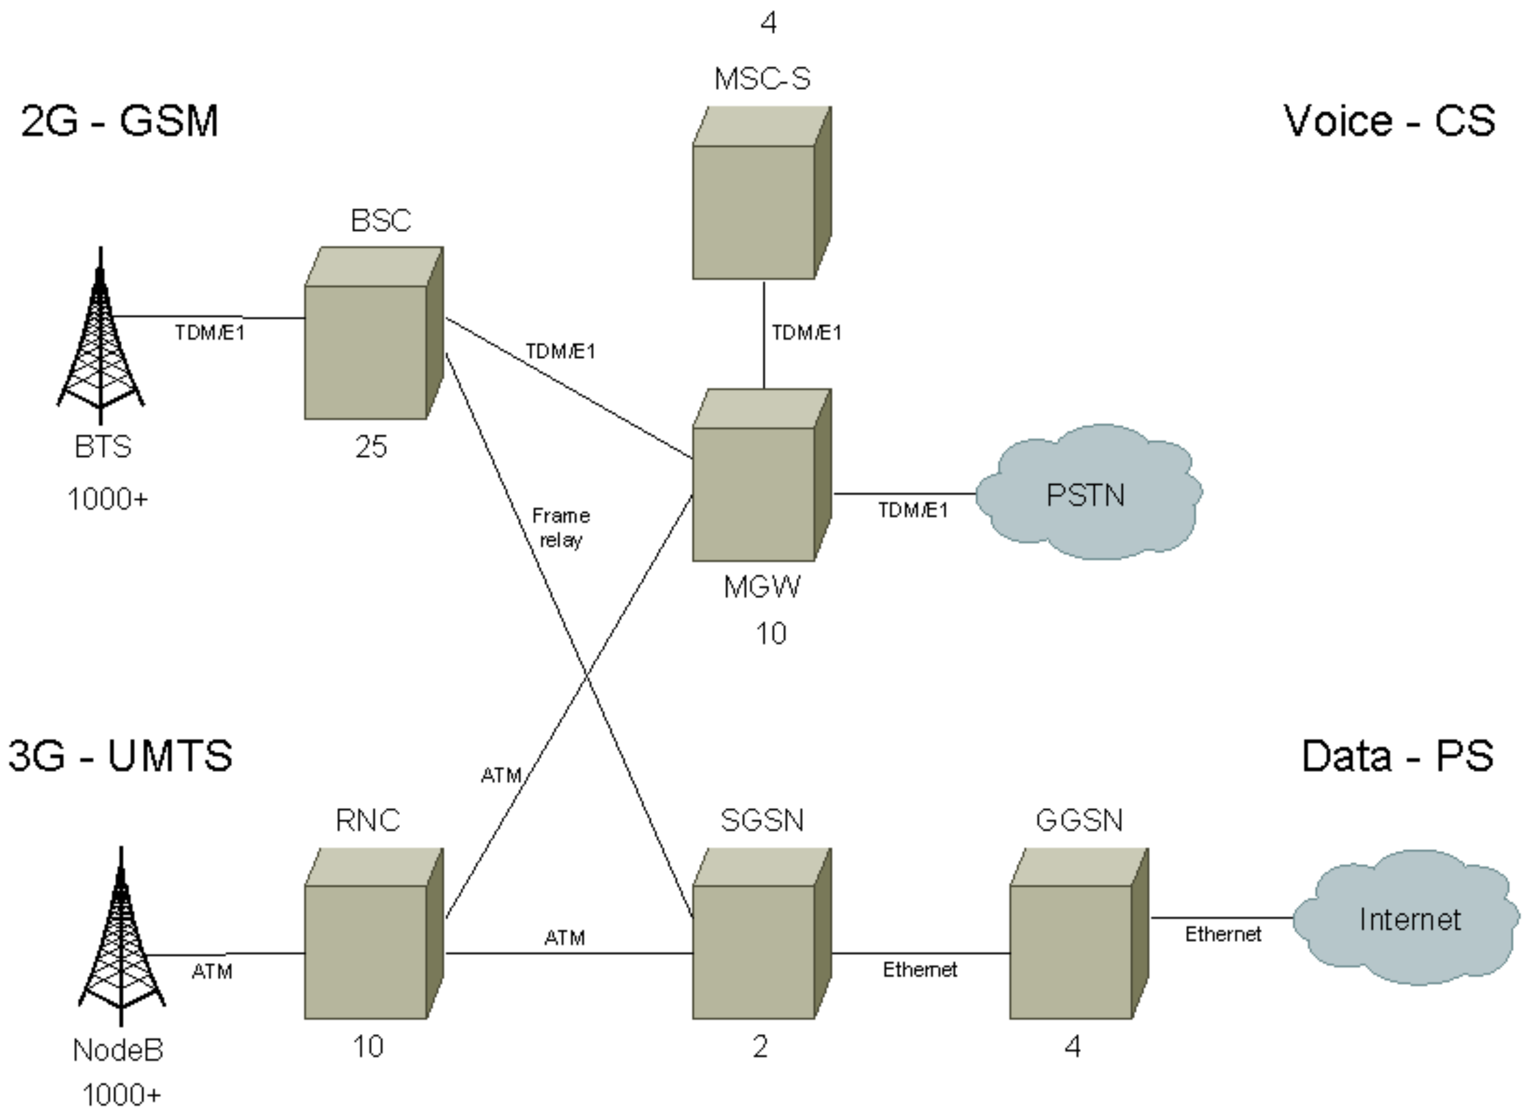

State of IP/ethernet migration oct. 2009

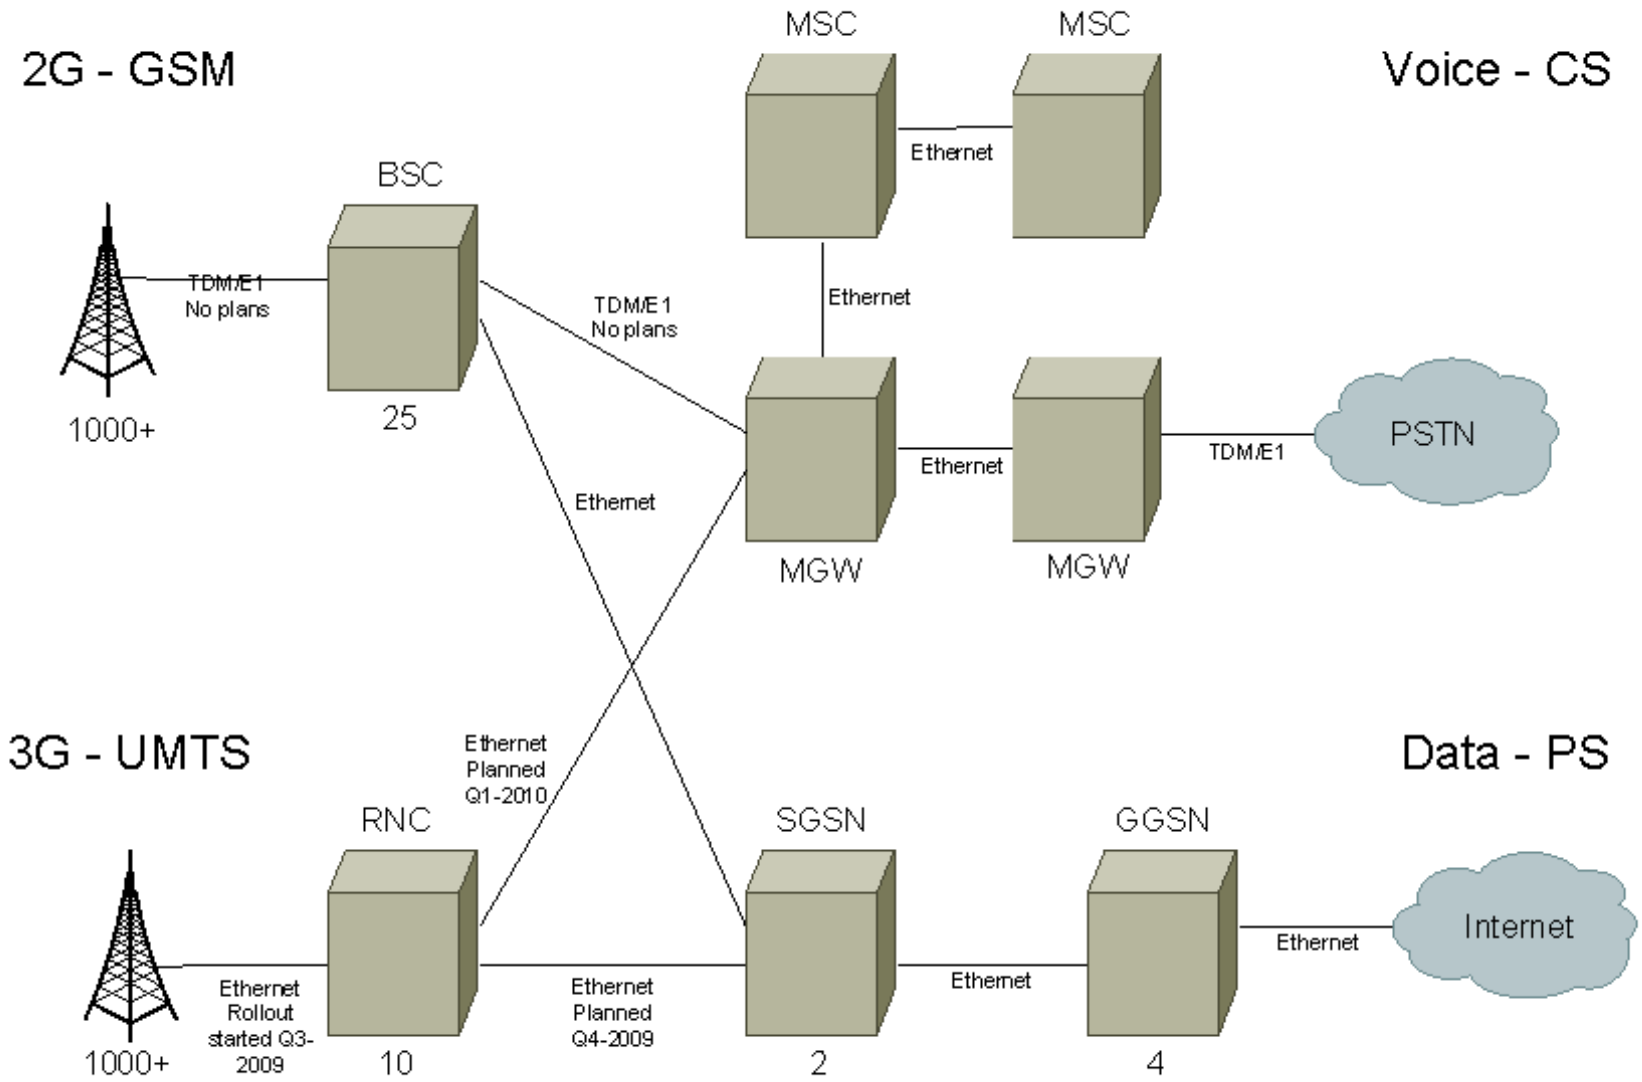

Core network requirements

- Very reliable (the five 9's)
- Sub second failover
- 3 QoS classes
- Handle high number of pps
- Low jitter
- Handle traffic growth (15%/year on voice – 100%+ on data)

What are our requirements

- Total 50.000 erlang – 15.000 on busiest site
- Room for growth
- Packet loss – < 0.00001%
- Operational in 3 months on 10 locations
- Main focus is to support Nb traffic
- Support IP RAN (Iub interface)
- Support other interfaces – IuCS, IuPS etc

Various scenarios

- Connect routers to existing circuits
- Build dedicated network
- Use existing network
- Use existing core and distribution

Connect routers to existing circuits

- Use existing circuits and connect routers using ATM
- Expensive
- Hub’n’spoke design
- No leadtime on circuits
- Dedicated equipment
- No constrains on software versions, config etc
- Support enterprise routing protocols (OSPF)
- Can use purpose selected equipment

(I’m sure some operators really do this!)

Build dedicated network

- Expensive
- Dedicated equipment
- Can use purpose selected equipment
- Use transport of choice – WDM, ethernet, POS
- No constrains on software versions, config etc
- Support enterprise routing protocols (OSPF)
- Can design for optimal performance

- Typical greenfield scenario

Our classic access network

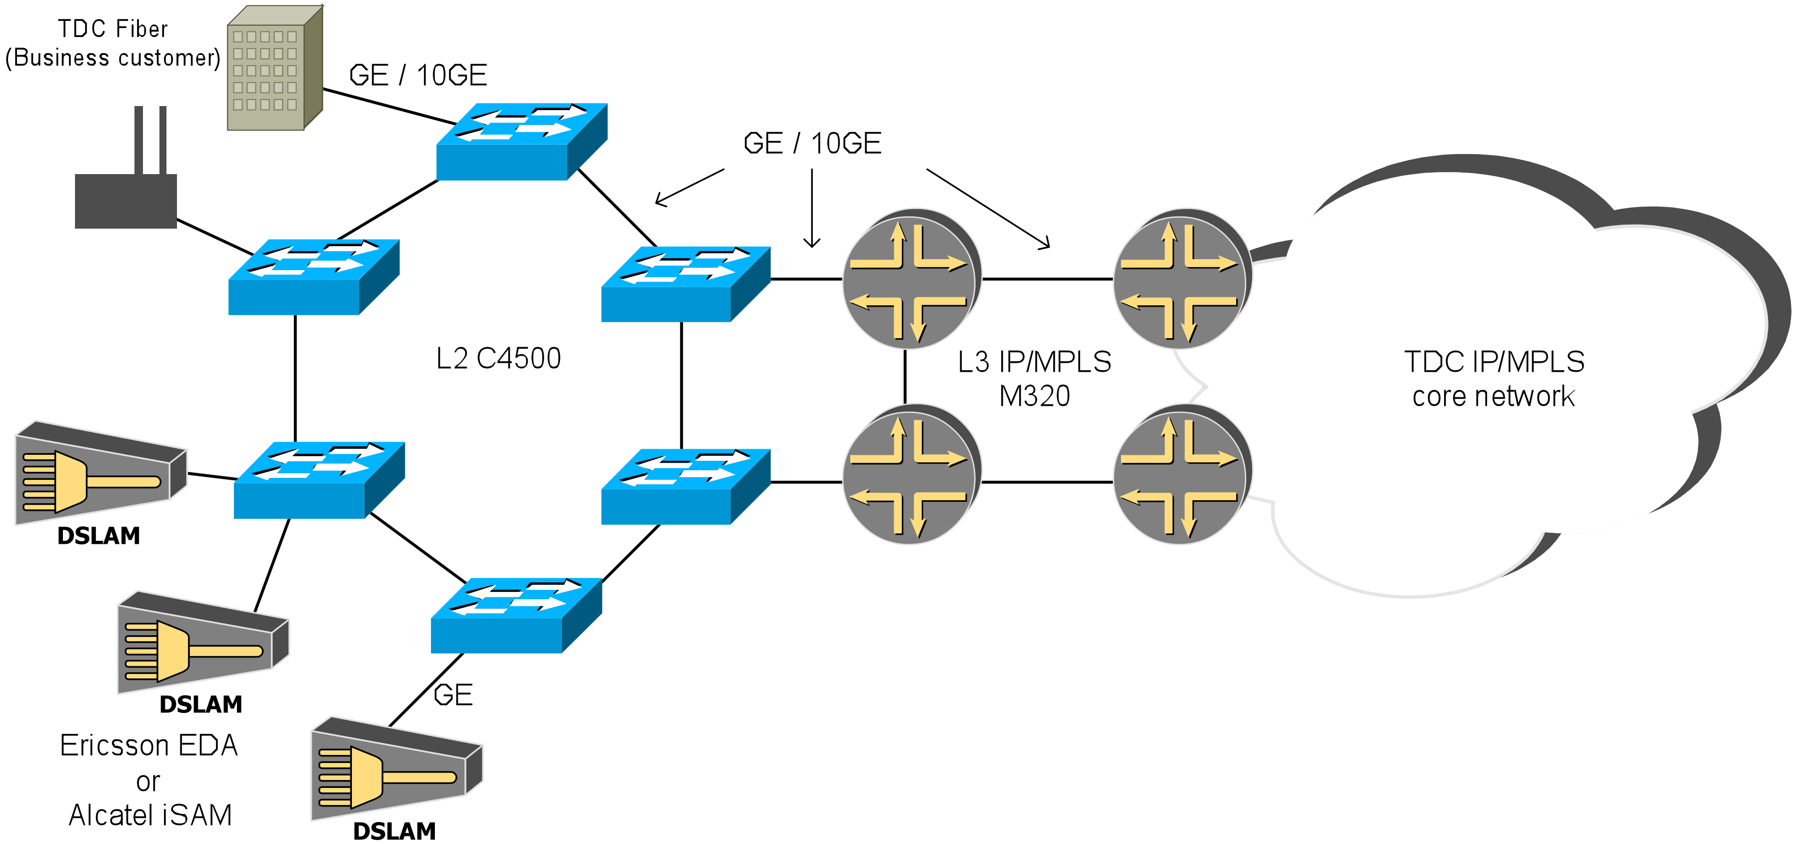

Use existing network

- Cheap
- Network already there
- Can be provisioned using existing tools
- Lot of daily changes
- Not possible to introduce mobile specific changes (like FRR, TE, newer software)
- Lots of customers – more prone to errors
- Broadcast in L2 rings may affect large number of mobile customers

(This is what we're using for IP RAN)

Use existing core and distribution

- Re-use expensive boxes
- Core and distribution is stable
- IP core sites and mobile cores sites are co-located
- Introduce dedicated PE’s for mobile PoP’s
- Allow us to use features like FRR, TE etc

We choose the last one!

- Reuse existing distribution – ~~GSR and~~ M320
- Reuse M10i's which were no longer in service
- No special config in existing network
- Since most routers and interface was in stock – cheap and fast deployment
- Roll-out of new mobile core done in approx one month

Design

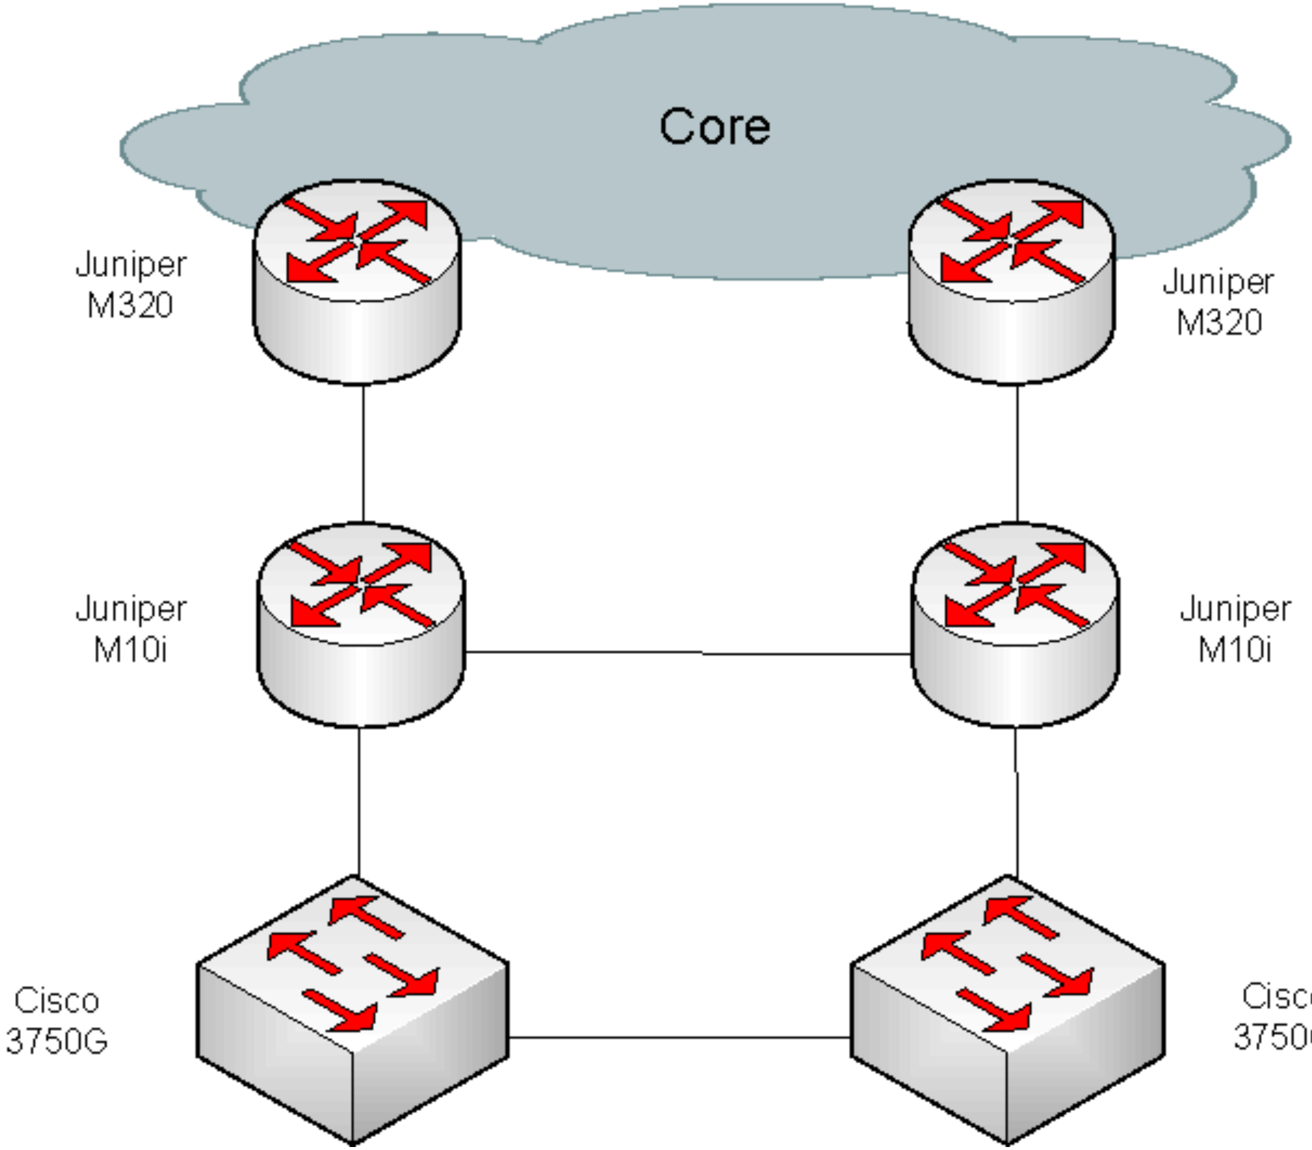
Core
Juniper
M320
Juniper
M320
Juniper
M10i
Juniper
M10i
Cisco
3750G
Cisco
3750G

Things learned

- Dedicated routers on the edge was good decision
- Fast rollout. From first meeting to deployment – 3 months
- Core and distribution has been really stable
- Uncompressed vs compressed voice makes a huge difference
- M10i doesn't do sub-second failover with full customer routes
- Our backbone QoS profiles wasn't good for voice
- Juniper does default ingress QoS on 10G IQ2 cards
- Size of IP packet overhead does matter!

Fast failover time (re-routing)

- M10i with full routing does not give sub second failover – several seconds
- Removing Internet routes and use default helps a lot – 1-2 seconds
- Fast reroute via backup tunnel gives us sub-second – hundreds of ms
- At that time Junipers local repair feature was not public, but we tested a special build and it gives similar failover times, without having to configure backup tunnel
- Caveat – remember to reconfigure the backup tunnel destination when you replace your LER/distribution routers!

MPLS backup tunnel

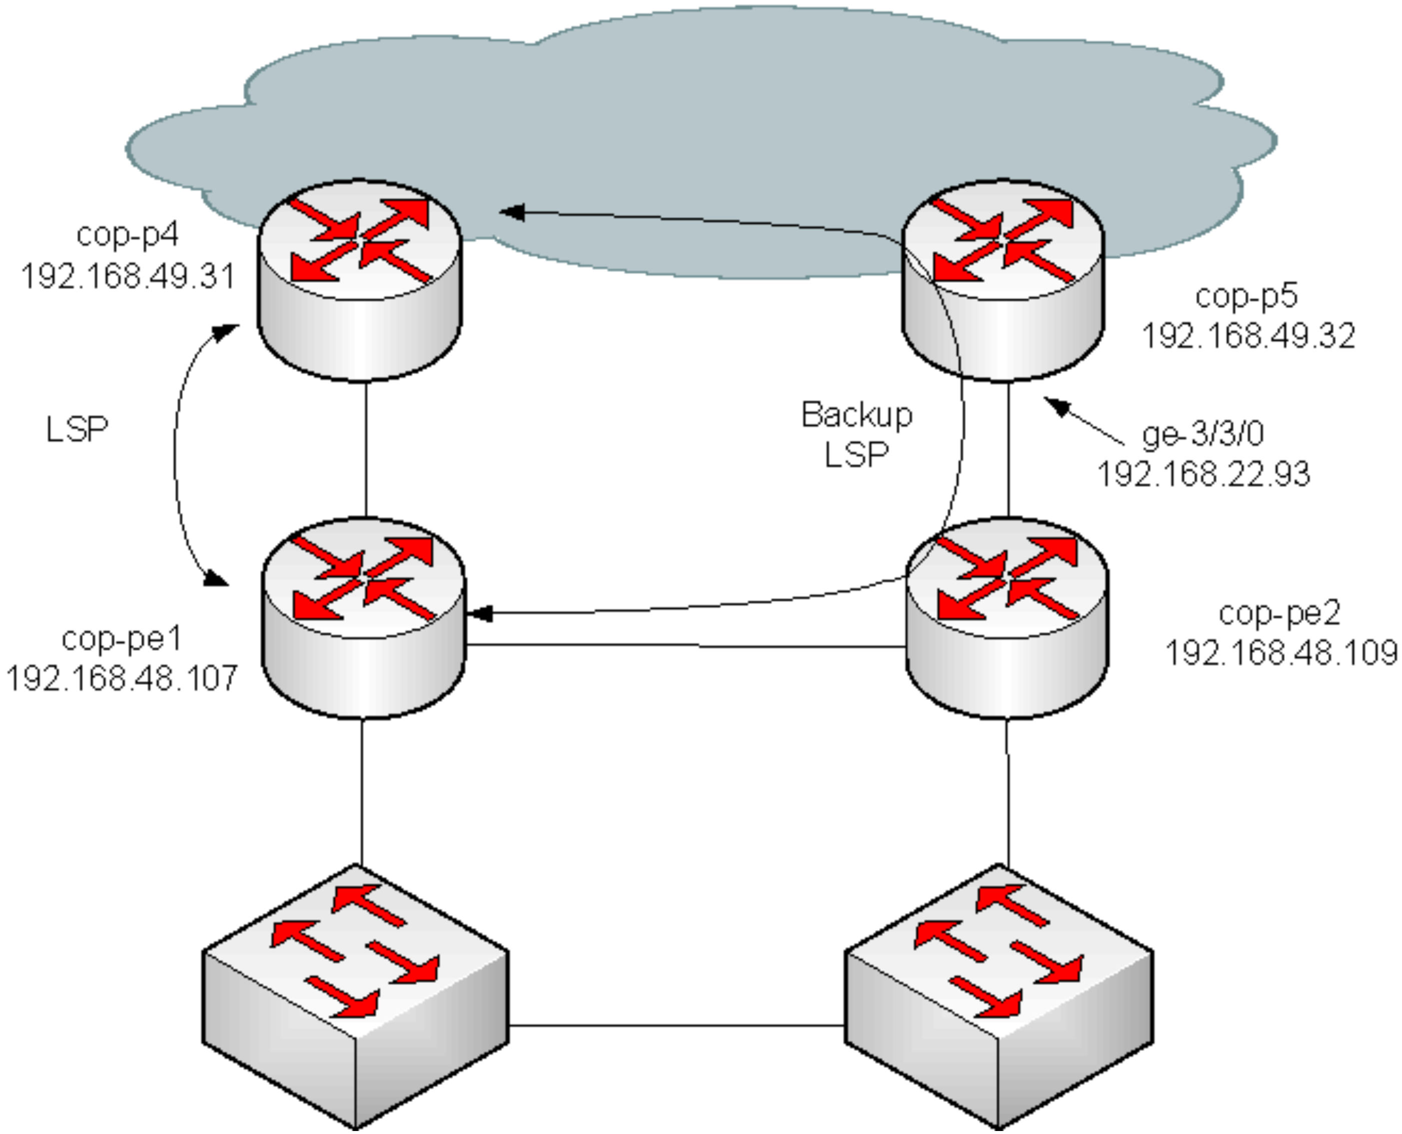

Fast re-route with backup LSP

```
chrille@cop-pe1> show configuration protocols mpls
...
label-switched-path coppe1-copp4-LP {
    to 192.168.49.31;                              // link to protect – not LSP target
    optimize-timer 60;
    description "T: coppe1-copp4 LP";
    link-protection;
}


chrille@cop-pe1> show mpls lsp
Ingress LSP: 1 sessions
To              From            State Rt ActivePath       P     LSPname
192.168.49.31   192.168.48.107  Up    734                 *     coppe1-copp4-LP
Total 1 displayed, Up 1, Down 0

Egress LSP: 0 sessions
Total 0 displayed, Up 0, Down 0

Transit LSP: 1 sessions
To              From            State    Rt Style Labelin Labelout LSPname
192.168.49.32   192.168.48.109  Up        0  1 SE  258288   248416 Bypass->192.168.22.93
Total 1 displayed, Up 1, Down 0
```

Backbone QoS should handle lot of EF

- Several issues with our QoS config surfaced when we added live traffic
- Nb (voice) – EF
- SS7 – AFnb
- At least 2/3 traffic is voice

- How much EF traffic do you have on your backbone links?
- 15% for EF on a gigabit link carrying Nb traffic is not enough!
- We used 65% on GSR links – dedicated policy map

Utilization on typical link

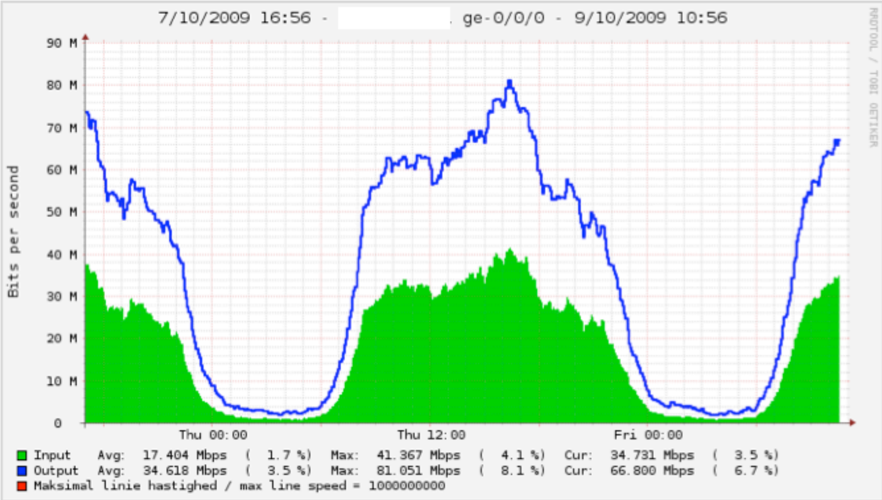

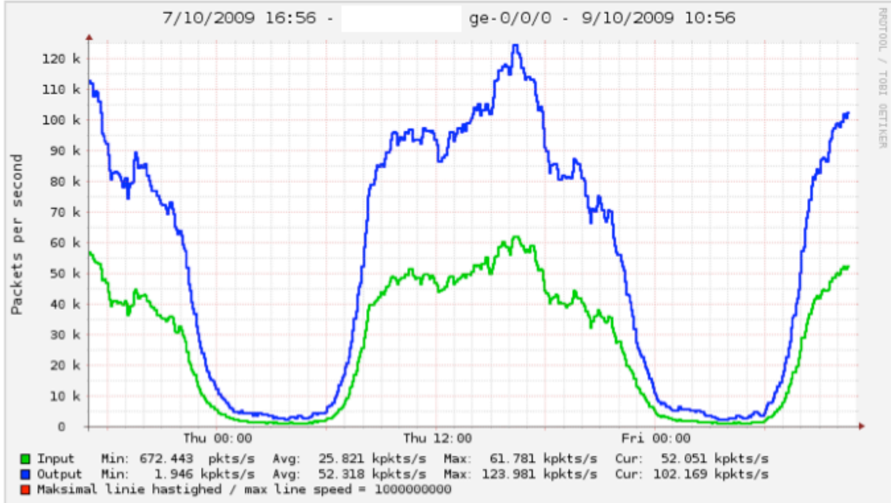

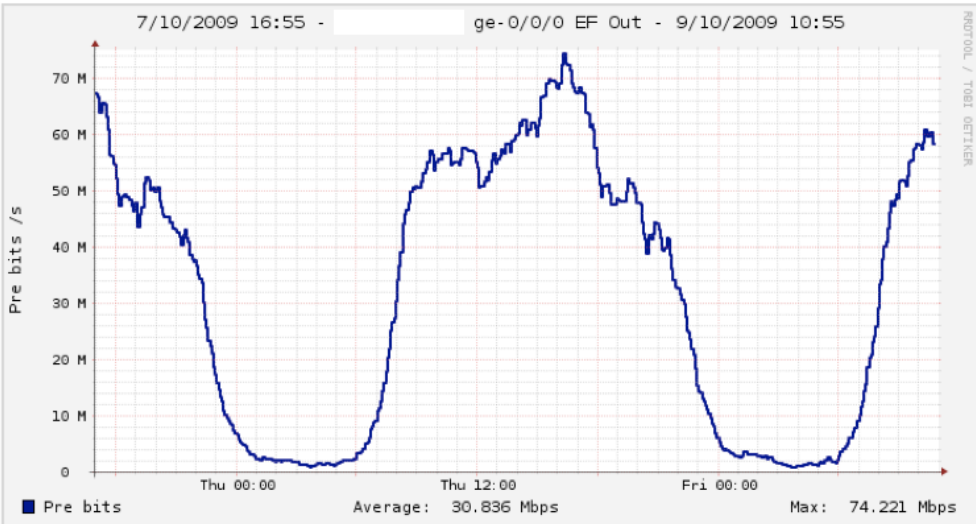

Impact of non-compressed voice

- We have been running PCM (un-compressed) speech

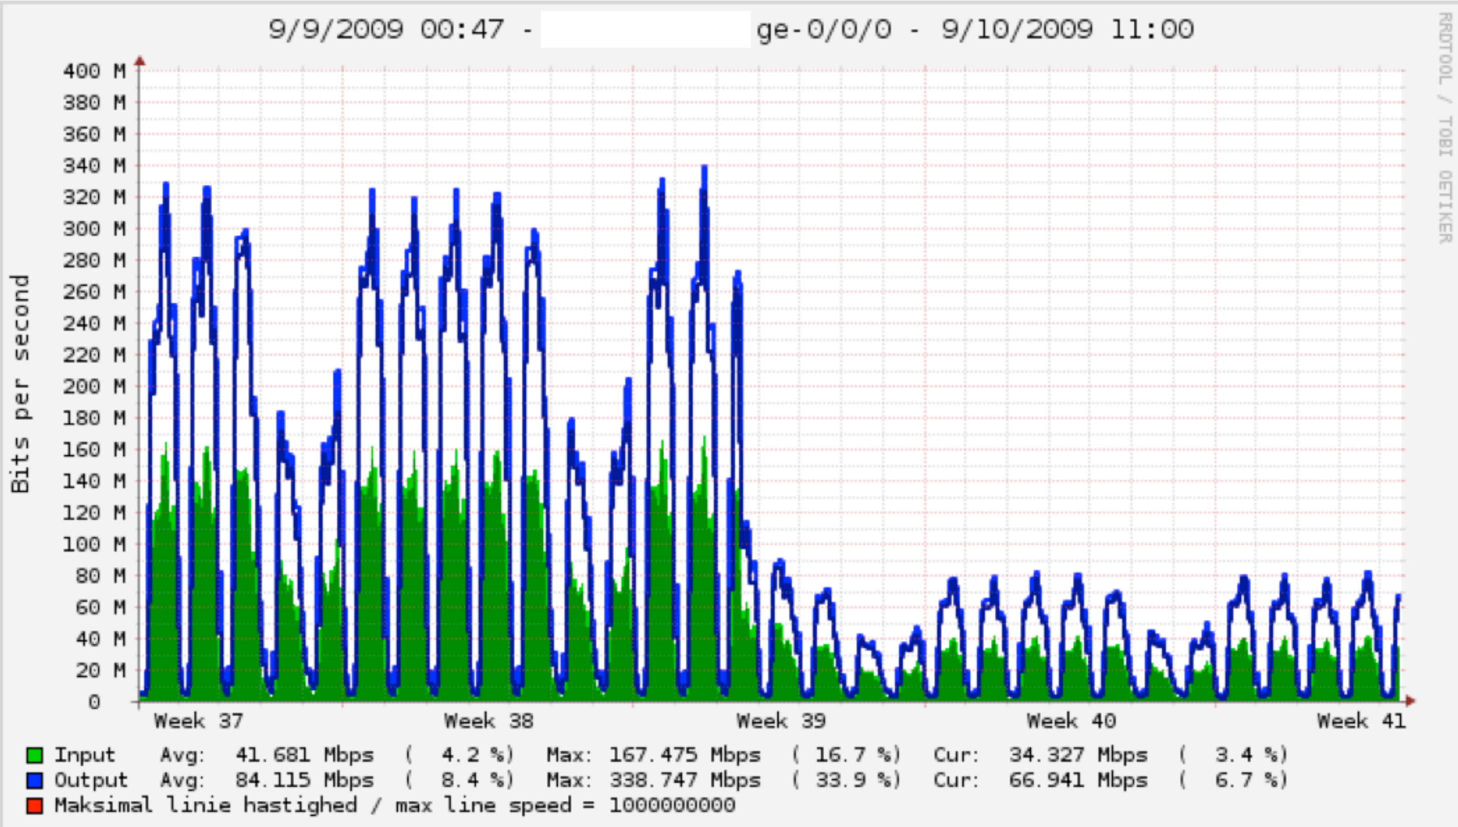

Default QoS settings is bad

- Juniper does default ingress QoS on 10G IQ2 cards
- Learned when PCM was introduced
- Default config sends EF traffik into 5% queue
- Disable default ingress queuing

10G IQ2 PIC default ingress QoS

```
chrille@labrouter> show class-of-service scheduler-map
Scheduler map: <default>, Index: 2

  Scheduler: <default-be>, Forwarding class: QUEUE-BE, Index: 19
    Transmit rate: 95 percent, Rate Limit: none, Buffer size: 95 percent,
    Priority: low
    Drop profiles:
      Loss priority    Protocol     Index     Name
      Low              non-TCP          1     <default-drop-profile>
      Low              TCP              1     <default-drop-profile>
      High             non-TCP          1     <default-drop-profile>
      High             TCP              1     <default-drop-profile>

  Scheduler: <default-nc>, Forwarding class: QUEUE-EF, Index: 21
    Transmit rate: 5 percent, Rate Limit: none, Buffer size: 5 percent,
    Priority: low
    Drop profiles:
      Loss priority    Protocol     Index     Name
      Low              non-TCP          1     <default-drop-profile>
      Low              TCP              1     <default-drop-profile>
      High             non-TCP          1     <default-drop-profile>
      High             TCP              1     <default-drop-profile>
```

…

10G IQ2 PIC default ingress QoS

```
chrille@labrouter> show interfaces xe-4/0/0 extensive
…
 Ingress queues: 8 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 QUEUE-BE                  136584               136584                    0
    1 QUEUE-AF_B                     4                    4                    0
    2 QUEUE-AF_NB                    4                    4                    0
    3 QUEUE-EF                    8281                 8281                    0
    4 QUEUE-NC                   54563                54563                    0
  Egress queues: 8 supported, 5 in use
  Queue counters:       Queued packets  Transmitted packets      Dropped packets
    0 QUEUE-BE                     680                  680                    0
    1 QUEUE-AF_B                     0                    0                    0
    2 QUEUE-AF_NB                    0                    0                    0
    3 QUEUE-EF              4294967296           4294967296                    0
    4 QUEUE-NC                  170371               170371                    0
```

- Disable

```
set chassis fpc 4 pic 0 traffic-manager mode egress-only
```

(or use apply group to disable on all interfaces)

Size does matter!

- Be aware of packet overhead!

IP Packet	AMR 12.2 codec
Payload size	31 bytes
Nb header	4 bytes
RTP header	12 bytes
UDP header	8 bytes
IP header	20 bytes
IP packet size, total	75 bytes

Ethernet overhead	Ethernet II
FCS	4 bytes
Ethernet frame	14 bytes
VLAN tag or MPLS labels	4 bytes
Preamble	8 bytes
Interframe gap	12 bytes
Ethernet framing, total	42 bytes

Ethernet overhead

For Nb: (75+42)/75 = 1.56 = 56% overhead!

For typical 1k packet: (1000+42)/1000 = 1.042 = 4.2% overhead

- We will saturate a gigabit link with 641 Mbps of payload traffic

- Different router vendors seems to count packet size different
- Juniper: Only counts L3 part of packet
- Cisco: Includes entire(?) ethernet frame

- Reason why we see different BW usage on Cisco L3, Juniper L3 and Cisco L2

What we would change

- One media VRF – one signaling VRF (+ Iub)
- Merge Nb, IuCS, Gb, IuPS, IuCS and Gn VRF
- Keep Iub in seperate VRF
- The road to LTE and direct tunnels
- Use combined L2/L3 device – eg MX, 7600

Current use of VRF's and 3G DT

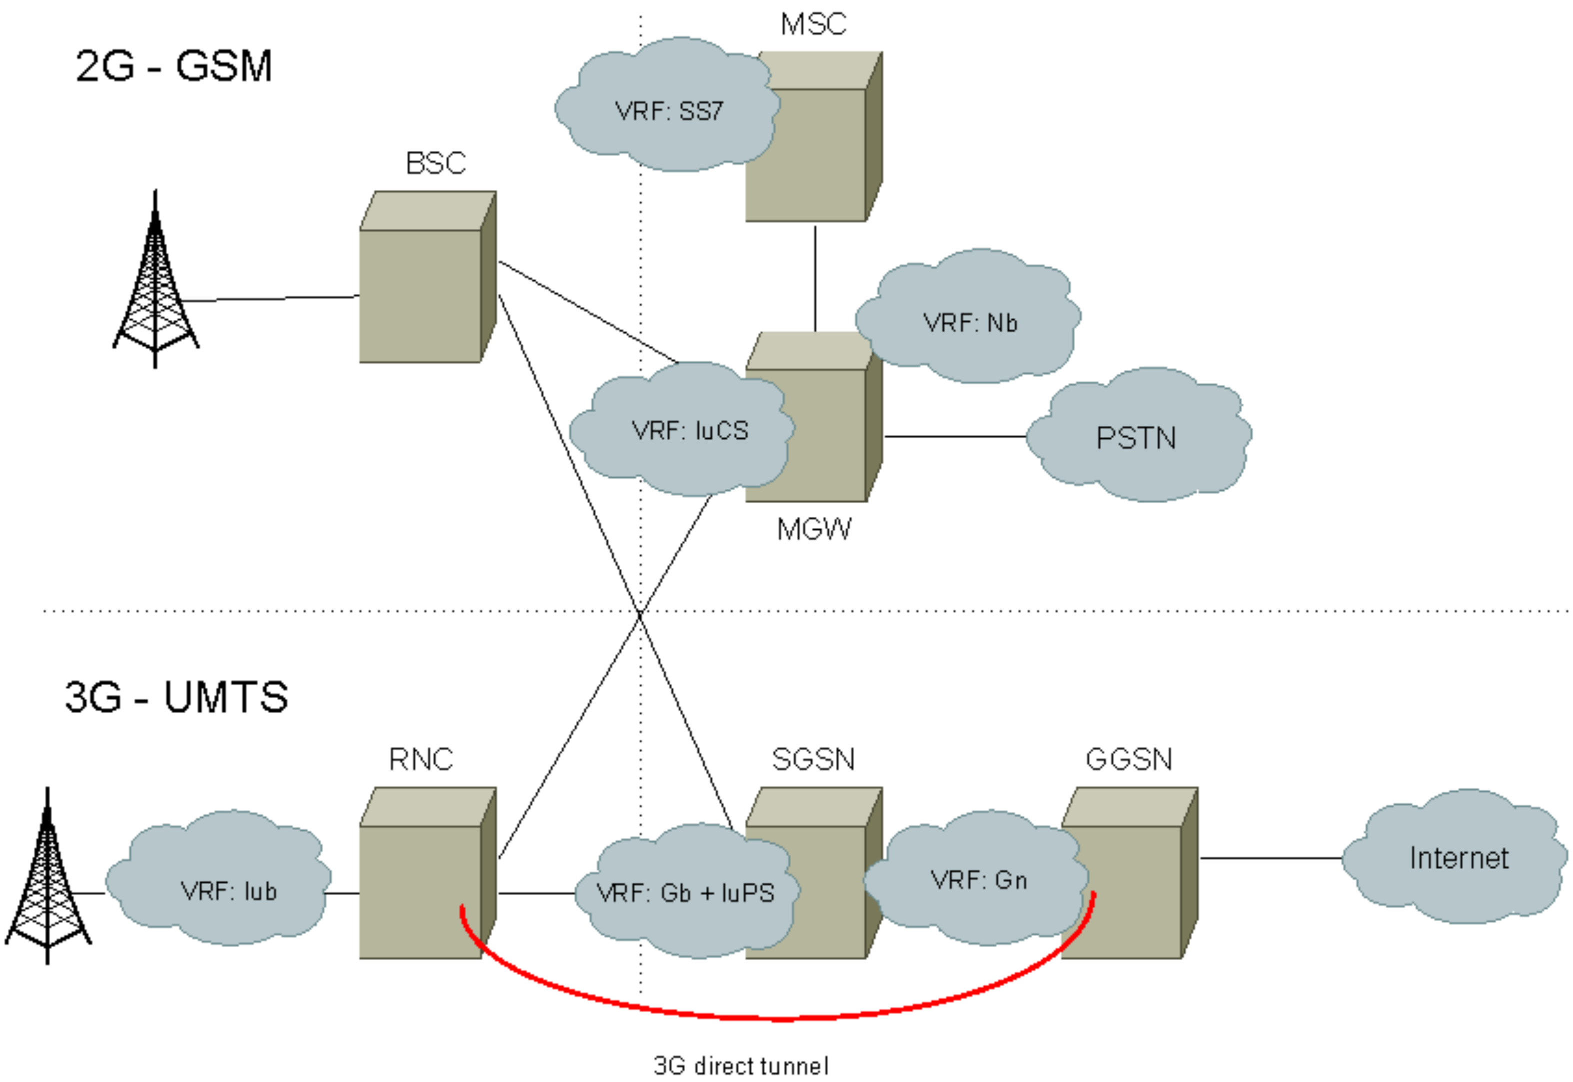

Dictionary

- 3G DT: 3G direct tunnel. Allows data userplane traffic to flow directly between RNC and GGSN – bypassing SGSN
- AMR: Adaptive Multi Rate. Audio compression codec widely used in GSM and UMTS networks
- BSC: Base station controller. Controls 2G basestations
- BTS: Base transceiver station. 2G basestation
- CS: Circuit switched – often meaning “voice”
- Erlang: Unit for measuring telephony load. One active call is one erlang
- GGSN: Gateway GPRS support node. Router between GPRS network and an IP net (Internet or VPN). Kinda like a BRAS
- GSM: Global System for Mobile communication. 2G network
- Gb interface: Interface between BSC and SGSN (2G data)
- Gn Interface: Interface between SGSN and GGSN
- IuCS interface: Interface between RNC and MGW (3G voice)
- IuPS interface: Interface between RNC and SGSN (3G data)
- Iub interface: Interface between NodeB and RNC (3G voice and data)
- LTE: Long term evolution – next generation mobile network with higher speeds. Based purely on IP transport and a flattened architecture

Dictionary

- MGW: Media gateway. Media transcoding, echo cancel and DTMF generation
- MSC: Mobile switching center. Handles all switching in 2G networks
- MSC-S: Mobile switching center server. Also referred to as MSS
- Nb: Interface between MGW nodes
- NodeB: 3G basestation
- PCM: Pulse code modulation. Un-compressed speech codec
- PS: Packet switched. Often meaning "data"
- PSTN: Public switched telephony network. "The" telephone network
- RAN: Radio access network – the network between the terminal and the core
- RNC: Radio network controller. Controls 3G basestations
- SGSN: Serving GPRS support node. Aggregates data connections. Kinda like a DSLAM
- SIGTRAN: SS7 over IP. Uses SCTP for reliable transport
- SS7: Signaling system #7. Signaling protocol used in telephony networks. Also carries SMS
- UMTS: Universal Mobile Telecommunications System. 3G

Questions