



**GoS in Policy Enforcement and
Management: Towards the Evolved
Packet Edge**

A White Paper from GoS Networks

Contents

Executive Summary	3
Introduction.....	5
Network Policy Control.....	6
Quality of Service.....	6
QoS in the NGN.....	7
Figure 1: PCRF & IMS Relationship in 3GPP Networks.....	8
Policy Frameworks in Fixed Networks	9
Figure 2: TISPAN RACS & NASS Framework.....	9
Limitations to Policy Enforcement.....	10
Figure 3: Policy Delivery to the CNG	11
Figure 4: HRAC Extensions	11
Policy Enforcement.....	12
Policy Reporting and Monitoring.....	13
Dynamic Policy Enforcement: Towards an Evolved Packet Edge	13
Figure 5: GoS in NGN Policy Architectures	14
Figure 6: The Effects of GoS.....	15
Conclusion	15
References.....	16
Abbreviations	16

Executive Summary

The recent explosion in broadband data consumption could aptly be described as a 'double-edged sword' for network operators. While this long-awaited boom in broadband data consumption is good news for operators, this rapid and strong growth in data traffic is also creating a 'capacity crunch', with networks of all types creaking under the strain of IP traffic demands. Enabled by flat-rate, all you can eat data plans, subscribers are eagerly consuming more and more services and capacity. The result is that subscribers are becoming increasingly aware of inconsistent quality levels for the services they consume, while service providers are becoming increasingly concerned over the strain that is being placed on the network.

The obvious solution to the issue of maintaining performance levels is to increase the capacity available for subscriber access. However, this is costly. Capacity increases are clearly essential, but the forecast growth is so great that despite every feasible upgrade, the network will remain as strained as it is now, if not more so.

Service providers are recognising that a more effective and intelligent solution to throwing capacity at the problem, is to take a proactive stance and actively manage access to and delivery of different services. Such an approach will enable service providers to attract premiums for guaranteed service performance and delivery, as well as ensure fair usage of existing plans and packages. This has brought the concepts of policy and QoS to the forefront of current debate. Policy creation, management and enforcement is now recognised as a significant revenue opportunity.

In order to deliver cost-effective, profitable network-based policies, a comprehensive framework based on the reliable delivery of QoS needs to be deployed. 3GPP has defined a key element, the PCRF, which has ultimate responsibility for the policies within a given network. Based on the policies in place, real-time QoS must be implemented via this series of additional network elements and functional entities.

Although originally defined for deployment in mobile networks, the PCRF solution has also been adopted for use in fixed networks, through parallel development of the RACS framework. However, when considering QoS, there are limitations to the current architecture,

Key Points

1. Policy will become a revenue generator, not just a revenue protector.
2. QoS is the means by which policies are enforced.
3. QoE demands reliable QoS; bandwidth is not the answer.
4. Successful policies, and hence the QoS framework, will address data demands at source i.e. the customer premises
5. GoS is the only solution that can reliably enforce policies for multiple, premium, real-time services per subscriber, while making efficient use of scarce access resources.
6. Deploying GoS to supplement the 3GPP / TISPAN policy framework as extended to CNGs will support the goal of monetising policies, and help future exploitation of data analytics information.

which negatively impact the ability of service providers to efficiently deliver QoS.

This TISPAN framework is gradually being extended directly to the customer premises. However, dynamic policy enforcement requires the deployment of QoS monitoring and enforcement software, directly into customer equipment. One solution to this requirement is offered by GoS Networks.

GoS Networks offers a unique solution called “*Guarantee of Service*”, which can be embedded as a software client in premises-based gateways. GoS allows the enforcement of differentiated policies that, in turn, become revenue generators.

GoS is the only solution that can reliably enforce QoS at the network edge, and hence allow for successful delivery of policies, for multiple, real-time services on a per subscriber basis, while making use of scarce access resources. Service providers can effectively monetise their policy framework by using GoS to support emerging TISPAN standards, directly to customer premises.

Introduction

The recent explosion in broadband data consumption could aptly be described as a ‘double-edged sword’ for network operators. The phenomenal success of the iPhone, the surging growth of smartphones, and increasing netbook usage are all contributing to a long-awaited boom in broadband data consumption, which is good news for operators after a decade of waiting for data traffic predictions to be realised. However, with good news comes bad news. This rapid and strong growth in data traffic is also creating a ‘capacity crunch’, with networks of all types creaking under the strain of IP traffic demands. Enabled by flat-rate, all you can eat data plans, subscribers are eagerly consuming more and more services and capacity. The result is that subscribers are becoming increasingly aware of inconsistent quality levels for the services they consume, while service providers are becoming increasingly concerned over the strain that is being placed on the network. With global IP traffic of all forms forecast to grow dramatically, which in turn will exacerbate current problems, service providers are being forced to take drastic steps to address subscriber satisfaction in order to prevent churn to rival networks.

The obvious solution to the issue of maintaining performance levels is to increase the capacity available for subscriber access. However, this is costly – particularly in the light of traffic growth forecasts, which suggest that capacity requirements will face a fast-moving target for some years to come. According to industry forecasts, global IP traffic will quadruple by 2014, reaching a level of 64 exabytes per month (Cisco, 2010). Mobile IP traffic will grow at the fastest rate, with an expected CAGR of 108%, while fixed IP traffic will constitute the largest proportion overall – contributing more than 94% of the total volume (Cisco, 2010). There are clearly significant challenges in both domains, but while much attention has been focused on the rapid growth of mobile IP traffic, the overwhelming problems will be found in fixed networks. Capacity increases are clearly essential, but the forecast growth is so great that despite every feasible upgrade, the network will remain as strained as it is now, if not more so.

Service providers are recognising that a more effective and intelligent solution to throwing capacity at the problem, is to take a proactive stance and actively manage access to and delivery of different services. Such an approach will enable service providers to prioritise certain kinds of traffic and to develop pricing strategies that allow them to attract premiums for guaranteed service performance and delivery, as well as ensure fair usage of existing plans and packages. This has brought the concepts of policy and Quality of Service (QoS) to the forefront of current debate.

In order to deliver cost-effective, profitable network-based policies, a comprehensive framework based on the reliable delivery of QoS needs to be deployed. Policies define the services and capabilities available to users; QoS ensures that these are delivered to the specified level. Such a comprehensive policy framework must also take into account visibility of customer equipment and be able to monitor and enforce policies at the level of individual subscribers. It is important to note that current architectures are limited in that they only address capacity reservation once traffic passes from the last mile into the transport network.

This white paper will discuss and address the issues and solutions for implementing successful policies in fixed networks, and illustrate how a unique approach to QoS provides the means to deliver fair and consistent Quality of Experience (QoE) to subscribers. In turn, this allows service providers to enforce and deliver a growing range of sophisticated and creative service policies.

Network Policy Control

In essence, policy defines a set of rules that govern the services to which individual users have access and the way in which they are entitled to use them. In their infancy, policies were simply concerned with ensuring that the appropriate charge was applied to a particular service. However, more sophisticated networks and surging application demands have led to recognition that policies can enable more powerful service offerings and richer billing possibilities.

On one level, policy is an abstract notion; yet on another, it has very real implications for users. This is because, ultimately, policy is translated into something tangible. When a subscriber uses a particular service, their perception of performance is termed 'QoE'. It is becoming clear that there is a willingness among some consumers to pay for enhanced QoE, particularly in certain areas such as online gaming.

Whereas QoE is entirely subjective – though it can be monitored through subscriber feedback – QoS can be measured objectively. Indeed, telecoms networks deploy complex underlying procedures to ensure that the right service attains the right level of performance at the right time – these procedures are collectively known as 'QoS techniques'.

Furthermore, the enforcement of policies is impossible without tight integration with QoS procedures; in their absence, user QoE cannot be guaranteed. So if policies are to emerge as a key element in creative pricing and service delivery strategies, a robust QoS framework is required to support them. While much of the attention in this emerging area has been on mobile networks, fixed networks face a similar, but much larger, problem.

Quality of Service

QoS in telecommunications and IP networks refers to “measurable and qualitative techniques that enable [service providers] to select, control, predict and measure the level of QoS and to guarantee predictable behaviour [for subscribers]” (Copeland, 2009). QoS is composed of a number of variable parameters that affect application performance. In IP networks, these include:

- Packet loss;
- Latency;

- Jitter;
- Corruption; and,
- Packet order¹.

In aggregate, these parameters constitute the final level of QoS.

Applications delivered across networks depend on the timely and accurate delivery of data, both supervisory and content related, irrespective of whether the service is a voice call, a multi-party video call, or a simple resource download. Each of these applications has different QoS requirements that must be met in order to deliver an acceptable QoE. Similarly, each application has different tolerance levels for degraded performance. For example, multi-party video sessions have a high requirement for real-time data because the threshold for service performance is relatively high as degradation below a certain point renders the session useless, whereas downloads of music files or large email attachments are activities that can be interrupted or treated with a lower priority.

In order to ensure that different applications are given the correct resources and can be delivered with the expected QoE (so, for example, a video chat session has an acceptable level of quality, and a voice call has appropriate audio quality), network-wide QoS solutions need to be deployed. This ensures that applications competing for bandwidth can be prioritised against their individual service requirements and therefore appropriate levels of service achieved. Without such QoS procedures, reliable policies are impossible unless unlimited, on-demand bandwidth is available.

Indeed, global migration towards converged, IP transport, core and service delivery networks has led to the formal inclusion of QoS frameworks within technical standards and greater recognition of the role that QoS has to play in the evolution of policy frameworks.

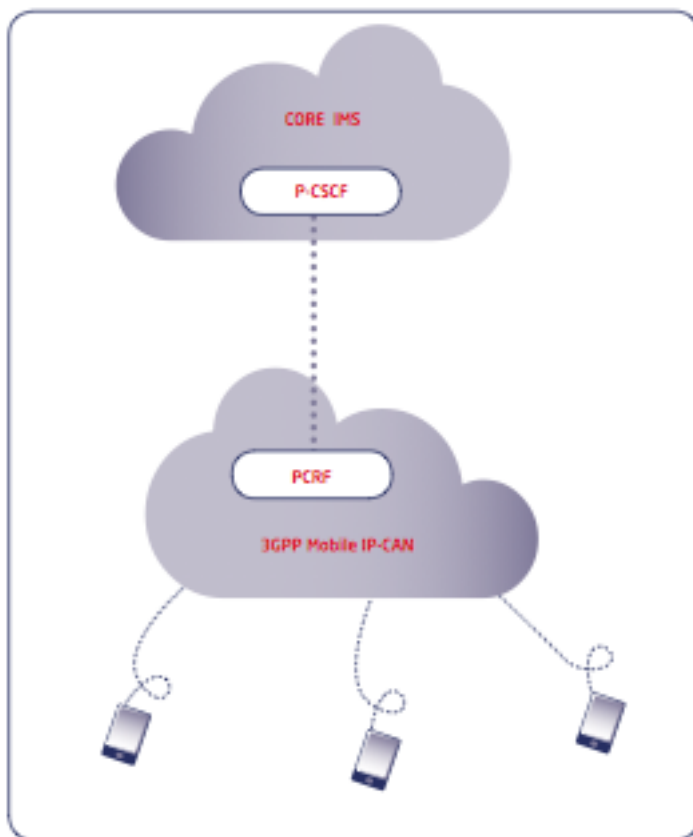
QoS in the NGN

Fixed and mobile networks have adopted the IP Multi-Media Subsystem (IMS) as the core engine for session control in NGN IP networks. This is tightly integrated with an extended framework for addressing QoS issues and ensuring that policies can be delivered to individual subscribers in real time. It is assumed that, in the future, connected user equipment (UE), such as mobile devices, laptops, softphones, and so on, will be IP-enabled and attach directly to what are known as IP Connectivity Access Networks (IP-CAN). These can be mobile, according to 3GPP standards; WIMAX or WLAN; or fixed via Ethernet and broadband DSL connections, according to TISPAN standards. In each case, the QoS required for an individual service supported by the network will be requested at session establishment and negotiated directly with the policy framework deployed in the core. If user requirements change within a session, renegotiation takes place, ensuring that the appropriate QoS adjustments are made.

¹ Note: depending on the behaviour of the receiver, corruption and out-of-order delivery may equate to additional packet loss.

3GPP has defined a key element, the Policy and Charging Rules Function (PCRF), which has ultimate responsibility for the policies within a given network. The PCRF is connected to additional, related functional entities that are associated with key network elements in order to complete the policy and charging architecture. Based on the policies in place, real-time QoS must be implemented via this series of additional network elements and functional entities. The PCRF framework can also be associated with the core session control platform – the IMS – allowing per-service based policies to be created, managed and enforced. A high level overview of this arrangement is provided in Figure 1.

Figure 1: PCRF & IMS Relationship in 3GPP Networks



Although originally defined for deployment in mobile networks, the PCRF solution has also been adopted for use in fixed networks. However, there are some additional considerations for fixed networks that have necessitated derivation of a slightly different architecture. This has been developed and specified under the work of the TISPAN group. TISPAN efforts have been integrated into the general 3GPP standardisation process.

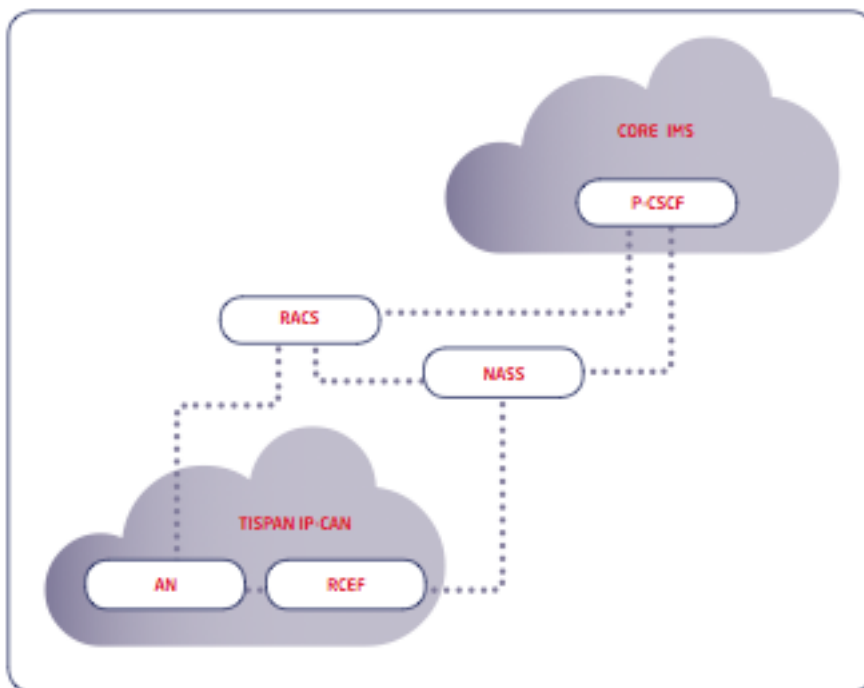
Policy Frameworks in Fixed Networks

TISPAN has defined a dedicated sub-system with responsibility for:

- Policy control;
- Resource reservation; and,
- Admission control.

The Resource Reservation and Admission Control Sub-system (RACS), provides “policy based control services to applications” (ETSI). The RACS allows applications to request the resources they need via standardised interfaces. As the RACS masks the underlying network, applications do not need to be aware of the network over which the service is to be delivered. Resources can be requested for real-time services and managed in real time to ensure that service upgrades or downgrades are achieved in line with user requirements. In addition to the RACS, there is a related subsystem, the Network Attachment Subsystem, which is responsible for the registration and initialisation of UE for access to TISPAN services. As such, it also has a role to play in the admittance of UE to a policy framework (ETSI). This architecture is illustrated in Figure 2.

Figure 2: TISPAN RACS & NASS Framework



Although the RACS and PCRF have emerged through the efforts of different, but related, standards bodies, the term PCRF is increasingly used as a generic term for a network policy control server. With the capabilities of the PCRF / RACS, service providers can deploy and maintain a single, all-encompassing policy framework across fixed, mobile and convergent networks.

The RACS can also interact with additional network elements that are supported in the transport network. However, when considering QoS, there are limitations to the current architecture, which negatively impact the ability of service providers to efficiently deliver QoS – and hence policies.

Limitations to Policy Enforcement

In order to accurately and cost-effectively provision network resources for individual applications and services (and therefore implement and enforce policies), service providers need to be able to differentiate between the different applications being used by individual subscribers. Whereas differentiation and policy enforcement can take place deeper in the network, this only really addresses the problem of competition for resources between subscribers, and not the competition between applications for the same subscriber. Because the access link is one of the principal bottlenecks, if applications using it are not directly under the control of the network policy framework, action is too late and other applications may be impaired. Policy monitoring and enforcement solutions need visibility of application and subscriber activity directly at the subscriber premises. This means that policy and QoS solutions also need to be applied directly to subscriber premises.

This problem with the current architecture has been recognised by TISPAN, and efforts have been made to address QoS and policy control within customer networks. In the latest generation of standards and research, TISPAN has now defined interfaces directly to Customer Premises Networks (CPN) and Customer Network Gateways (CNG) and has created a new entity, termed the H-RAC, or Home Resource and Admission Control entity (ETSI). This is shown in Figures 3 and 4.

Figure 3: Policy Delivery to the CNG

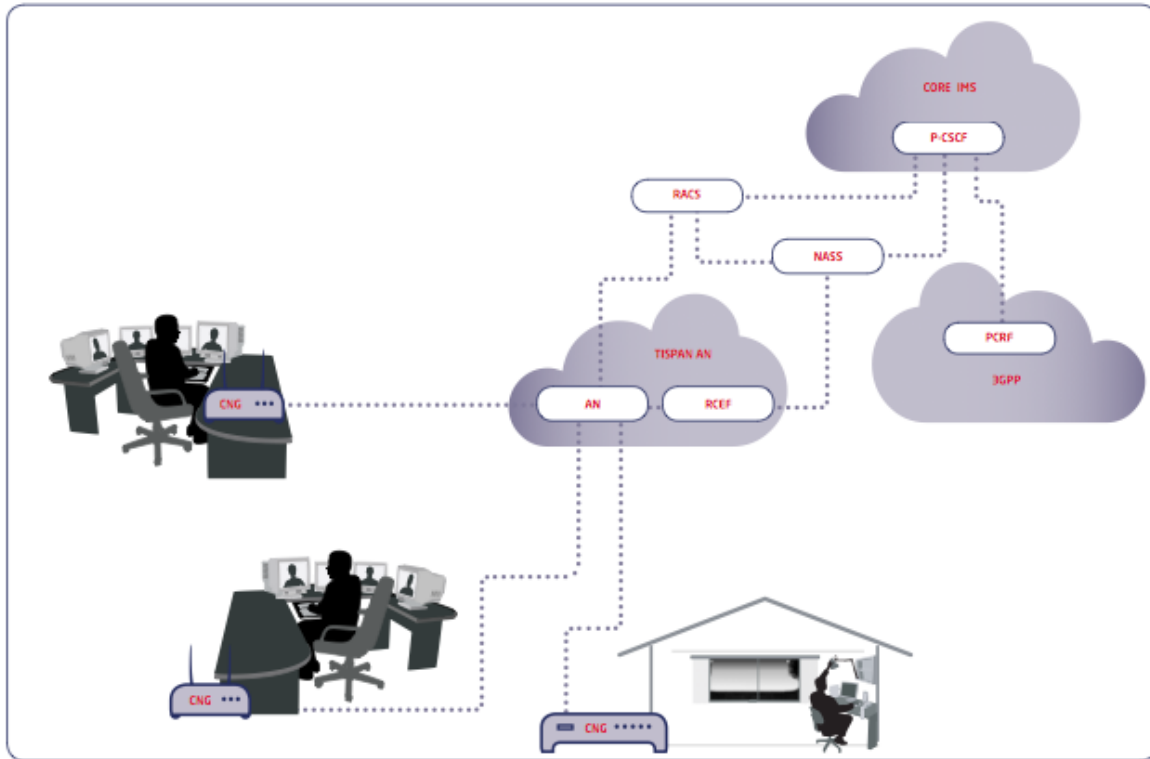
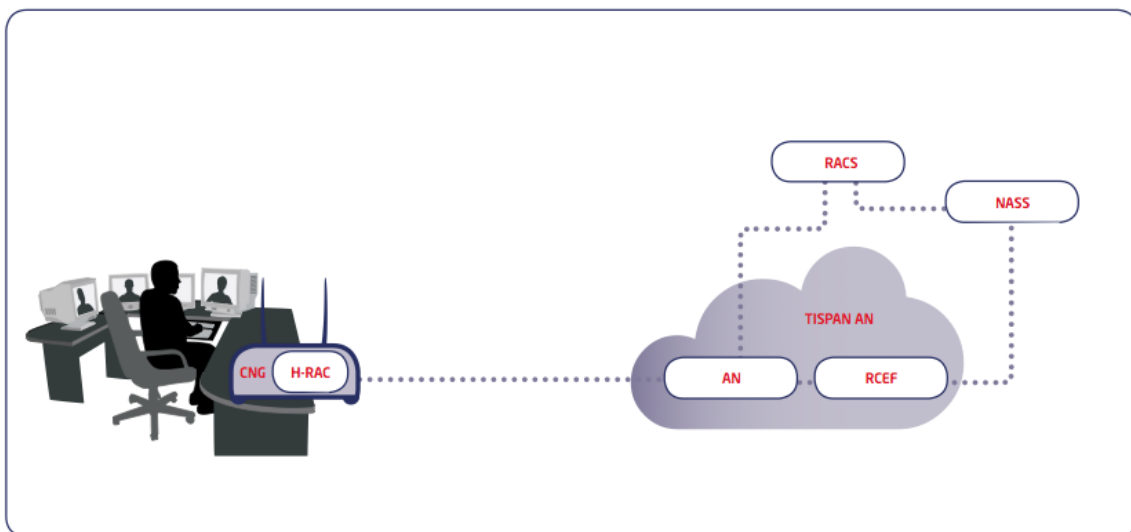


Figure 4: HRAC Extensions



The H-RAC is further decomposed into additional entities, beyond the scope of this paper, but the entity has to implement policy decisions in response to application demands locally, in conjunction with specialised software co-resident in the gateway. This software must be able to identify traffic, discriminate between applications, and enforce the required policy decisions. Furthermore, it must be able to interact with local QoS-capable software, so that the policy decisions can be enforced and translated into appropriate resource reservation and prioritisation.

Policy Enforcement

Decisions on QoS, determined by the agreed network policies can be extended directly to CNGs in fixed networks over broadband access. There may be a range of equipment connected to such CNGs, encompassing both consumer and business needs, including:

- Legacy analogue handsets;
- Legacy digital handsets;
- SIP extensions;
- Video terminals;
- PCs;
- Fax machines;
- Video on demand solutions;
- IPTV platforms; or
- Real-time gaming platforms.

Each of these devices is likely to have different needs, and the overall requirements of the users behind each CNG may change dynamically during each session. A CNG could be deployed in, for example, a home with multiple users, an enterprise, or an SME. By delivering QoS and policy control to the CNG, service providers can ensure that they enforce policies for subscriber traffic continuously, in real time, before problem traffic reaches the network transport layer. For example, if a residential subscriber is a heavy P2P user, a policy could be enforced, via their CNG, to reduce the bandwidth available for P2P applications in order to ensure capacity for other, real-time traffic, such as gaming or video chat. The problem can be eliminated before it even emerges, as the policy blueprint ensures that different traffic types are controlled at the point of origin, ensuring fairness to other users on the same network, or to other applications with real-time requirements for the same subscriber.

Policy Reporting and Monitoring

Similarly, the overall network policy framework has to adjust dynamically to changing requirements and demands. The local QoS function must be able to register such changing demands and allow the policy functions to manage these fluctuations in real time. Monitoring and reporting changing real-time requirements for user traffic at source thus enables dynamic policy control. For example, a service provider may adjust the amount of bandwidth available to P2P traffic based on a measurement of the overall network load, while local control would prevent ‘undesirable’ traffic from consuming network bandwidth, in core, access and backhaul domains. Similarly, the service provider could also permit P2P traffic to consume more bandwidth at the level of individual users once a real-time application has been terminated, but may wish to degrade the P2P traffic once again as demand resumes. By continuously monitoring and reporting traffic demands and usage via the overall policy framework, service providers can provide flexible and dynamic policies based on real-time user requirements.

Therefore, deployment of the appropriate QoS solution, embedded as a function within CNGs will enable service providers to offer multiple classes of service, rather than simply apply generic throttling techniques to traffic that originates from subscriber premises.

Dynamic Policy Enforcement: Towards an Evolved Packet Edge

In core networks, IP session routing, internetwork connectivity and service control is offered by the IMS. An adjunct policy framework allows network-based policy and charging rules to be implemented for all users, irrespective of their mode of access. This framework is gradually being extended directly to the customer premises, via the efforts of TISpan and the definition of new functions and capabilities. However, dynamic policy enforcement requires the deployment of QoS monitoring and enforcement software directly into the CNG. One solution to this requirement is offered by GoS Networks.

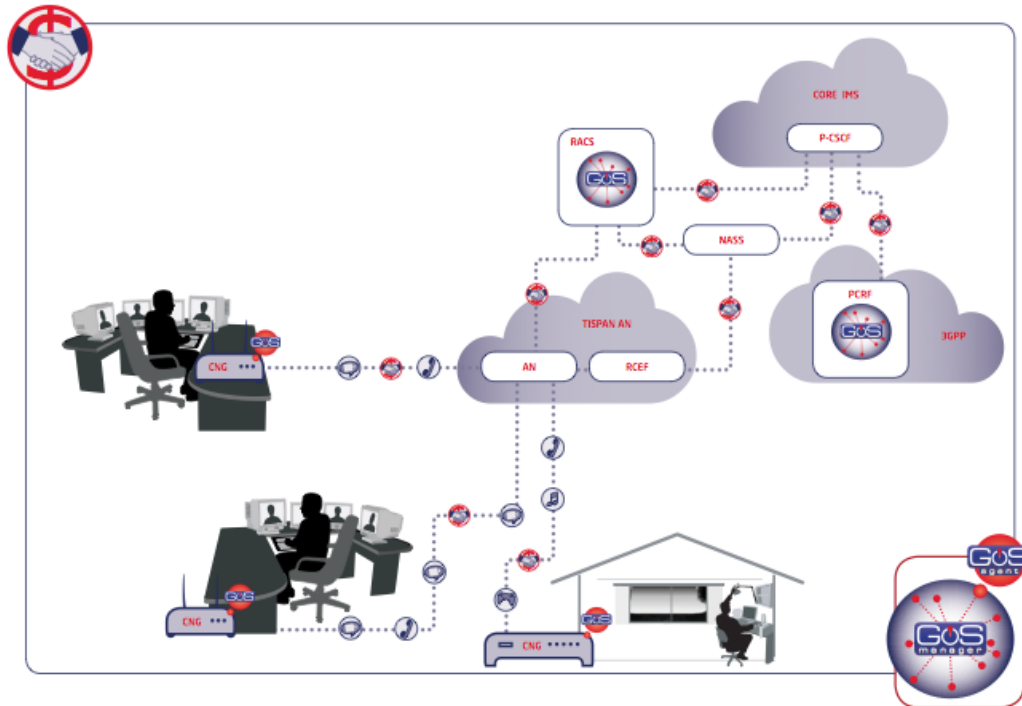
GoS Networks offers a unique solution called “*Guarantee of Service*”, which can be embedded as a software client in CNGs. GoS Agent provides the means for service providers, via a standardised policy control framework, to enforce policies directly at the traffic source – that is, at the extreme edges of the network. GoS Agent delivers multiple, real-time classes of service, enabling independent control of traffic for multiple, parallel real-time and near real-time applications:

- Throughput;
- Loss; and,
- Delay / Jitter.

GoS allows the enforcement of differentiated policies that, in turn, become revenue generators. Furthermore, GoS Agent is complemented by GoS Manager, which together provide a comprehensive


reporting solution that delivers real-time updates of traffic usage and demand to the policy engine, allowing dynamic policy enforcement in response to changing usage patterns. The relationship of GoS to the overall policy framework is illustrated in Figure 5.

Figure 5: GoS in NGN Policy Architectures



GoS goes beyond basic traffic policing approaches and simple call-admission control to allow delivery of a real-time, dynamic policy framework, supporting multiple classes of service and QoS control to ensure reliable QoE for all users. This evolved packet edge will become a key component in service providers' evolving policy and charging strategies, ensuring that they can maximise network efficiency and offer genuinely differentiated policies. An illustration of the effects of deploying GoS is shown in Figure 6.

Figure 6: The Effects of GoS



The monitoring and reporting capabilities of GoS also allow service providers to more effectively collect and analyse information on subscriber activities and application usage, which in turn can be used proactively to develop new policies and profiles, and enhance network optimisation. Data analytics is emerging as a key resource for service providers seeking to find new ways to anticipate and respond to service and bandwidth utilisation, and GoS provides an elegant solution.

Conclusion

Policy, as enabled and enforced in telecoms networks, is set to become a key revenue generator for service providers, as innovative new pricing plans and offers are introduced to enterprise and consumer customers alike. QoS is critical to the successful introduction of policies, as it provides the means by which the policies can be measured and enforced. Customers with a high QoE are happy customers, and so will be less likely to seek out competing services and alternative service providers.

Maintaining high levels of QoE requires reliable QoS. Although one approach is simply to increase available bandwidth, this is costly and unsustainable. An alternative, more effective approach is to implement network-wide QoS measurement and enforcement solutions. Although much effort has been expended in this area, the next issue to address will be monitoring and reporting data demands at the customer premises and, indeed, TISPAN is now attempting to do this via extension of policy mechanisms

to customer network gateways. When such solutions are realised, they will need specialised software to deliver the required QoS.

GoS is the only solution that can reliably deliver QoS, and hence allow for successful delivery of policies, for multiple, real-time services on a per subscriber basis, while making use of scarce access resources. Service providers can effectively monetise their policy framework by using GoS to support emerging TISPAN standards, directly to customer premises.

References

Cisco. (2010). *Visual Networking Index*.

Copeland, R. (2009). *Converging NGN Wireline and Mobile 3G Networks with IMS*.

ETSI. *TISPAN NGN Functional Architecture; Network Attachment Subsystem (NASS)*.

ETSI. *TISPAN Remote CPN QoS Control; Study on CPN - RACS Interaction*.

ETSI. *TISPAN Resource and Admission Control Subsystem (RACS); Functional Architecture*.

Abbreviations

3GPP *Third Generation Partnership Project*

CAGR *Compound Annual Growth Rate*

CNG *Customer Network Gateway*

CPN *Customer Premises Network*

EPC *Evolved Packet Core*

H-RAC *Home Resource and Admission Control entity*

IMS *IP Multimedia Subsystem*

IP *Internet Protocol*

IP-CAN *IP Connectivity Access Network*

NASS *Network Attachment Subsystem*

NGN *Next Generation Network*

P2P *Peer to Peer*

PCC *Policy and Charging Control*

QoE Quality of Experience

QoS Quality of Service

RACS Resource and Admission Control Subsystem

TISPAN Telecommunications and Internet converged Services and Protocols for Advanced Networking

UE User Equipment

WLAN Wireless Local Area Network