QoS: (still) a grand challenge?

David Hutchison
Lancaster University

Wikipedia definition:

"In the field of computer networking and other packet-switched telecommunication networks, the traffic engineering term quality of service (QoS) refers to resource reservation control mechanisms rather than the achieved service quality. Quality of service is the ability to provide different priority to different applications, users, or data flows, or to guarantee a certain level of performance to a data flow."

Cisco Internetworking Technology Handbook:

"QoS refers to the capability of a network to provide better service to selected network traffic over various technologies, including FR, ATM, Ethernet, 802.1, SONET and IP-routed networks that may use any or all of these technologies."

ITU-T QoS Framework:

"A set of qualities related to the collective behaviour of one or more objects."

Was QoS ever a grand challenge?

Yes, it probably was … in the early days …

0. ATM as universal carrier (B-ISDN)

1. How to get packet networks (ATM)
to cope with digital audio and video

2. How to get packet networks
to treat all traffic with its appropriate
level of service (ie as the end-user requires)

The early 1990s:

OSI still around; ‘coloured book’ protocols
Internet still not widely accepted; ATM
And the ‘new environment’ of DAV

There were several efforts to
produce a QoS architecture

Concurrently, traffic engineering
was the subject of much research

History (for me):

OSI'95 project (1990-93)

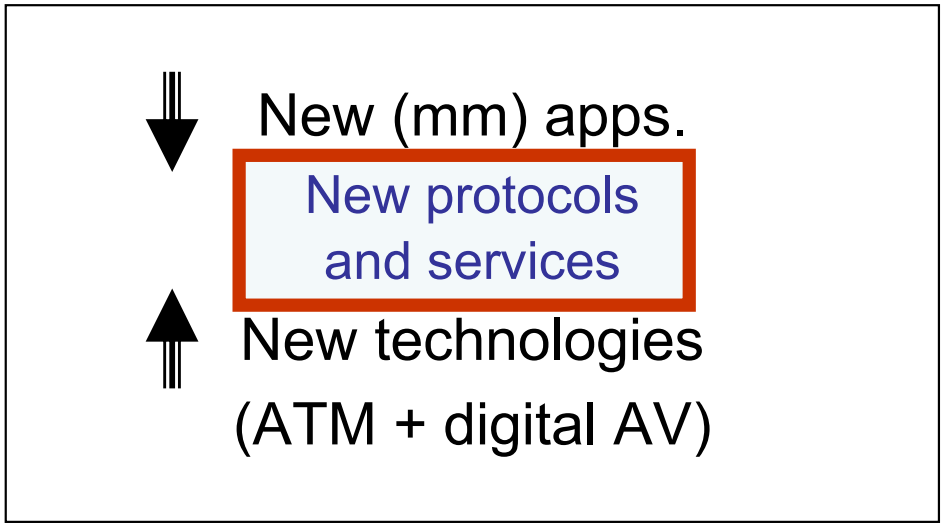

Note that:
IWQoS began in 1993, in Montreal

OSI'95 view of the world:

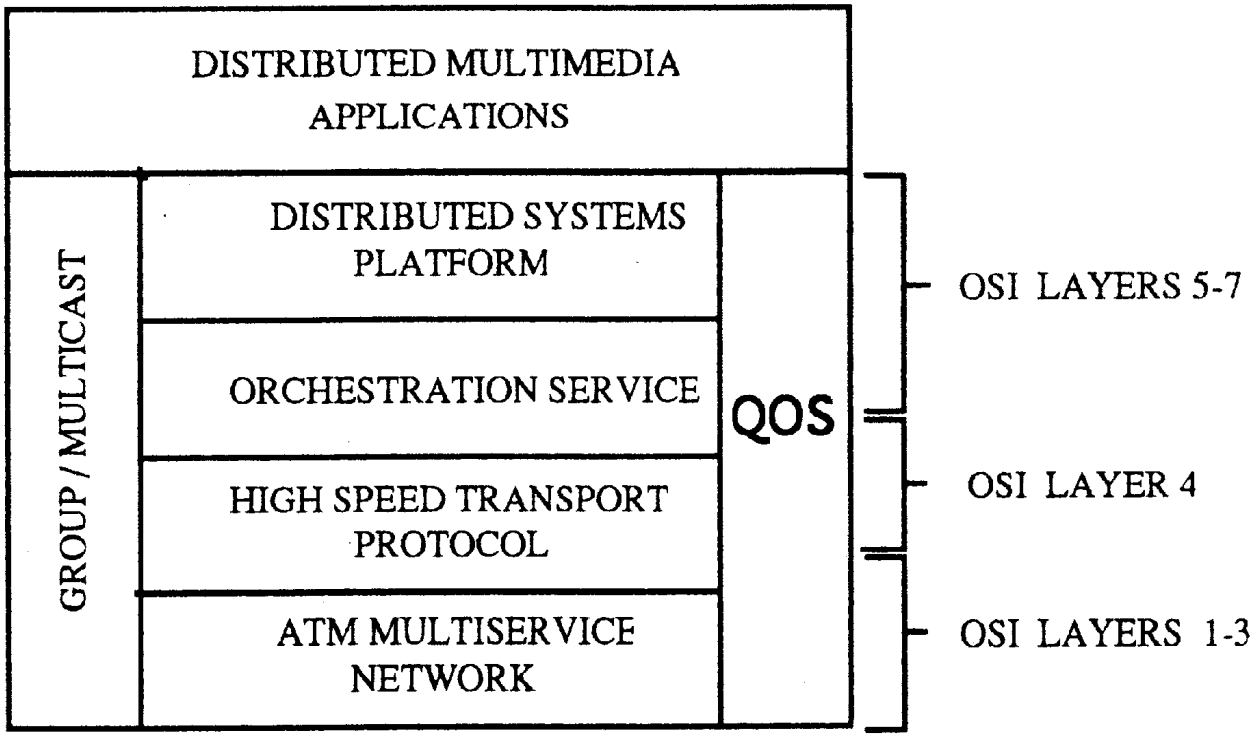

Fig. 11.1 Distributed multimedia architecture

An early attempt at QoS architecture:

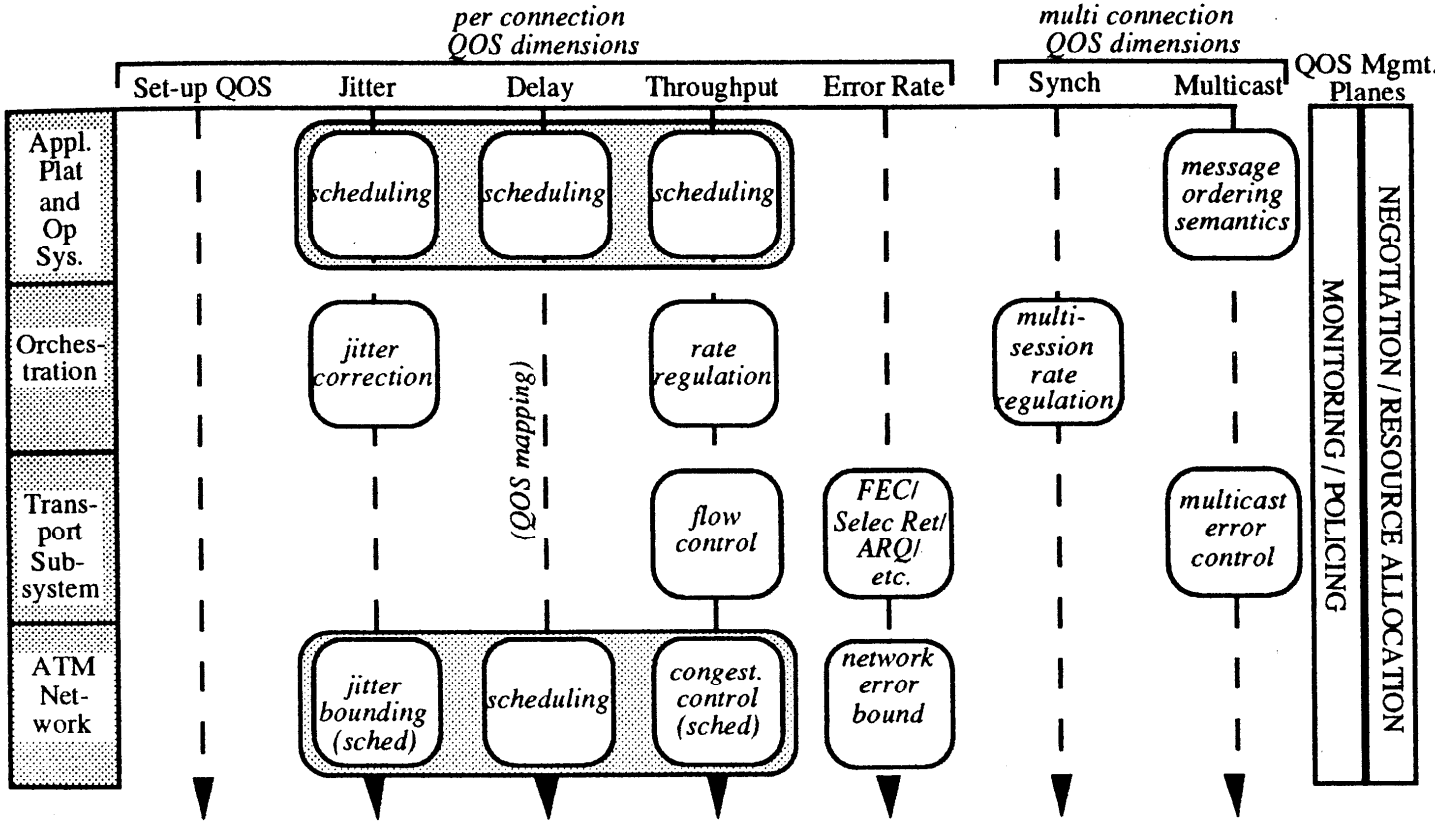

Fig. 4.2 QOS-A

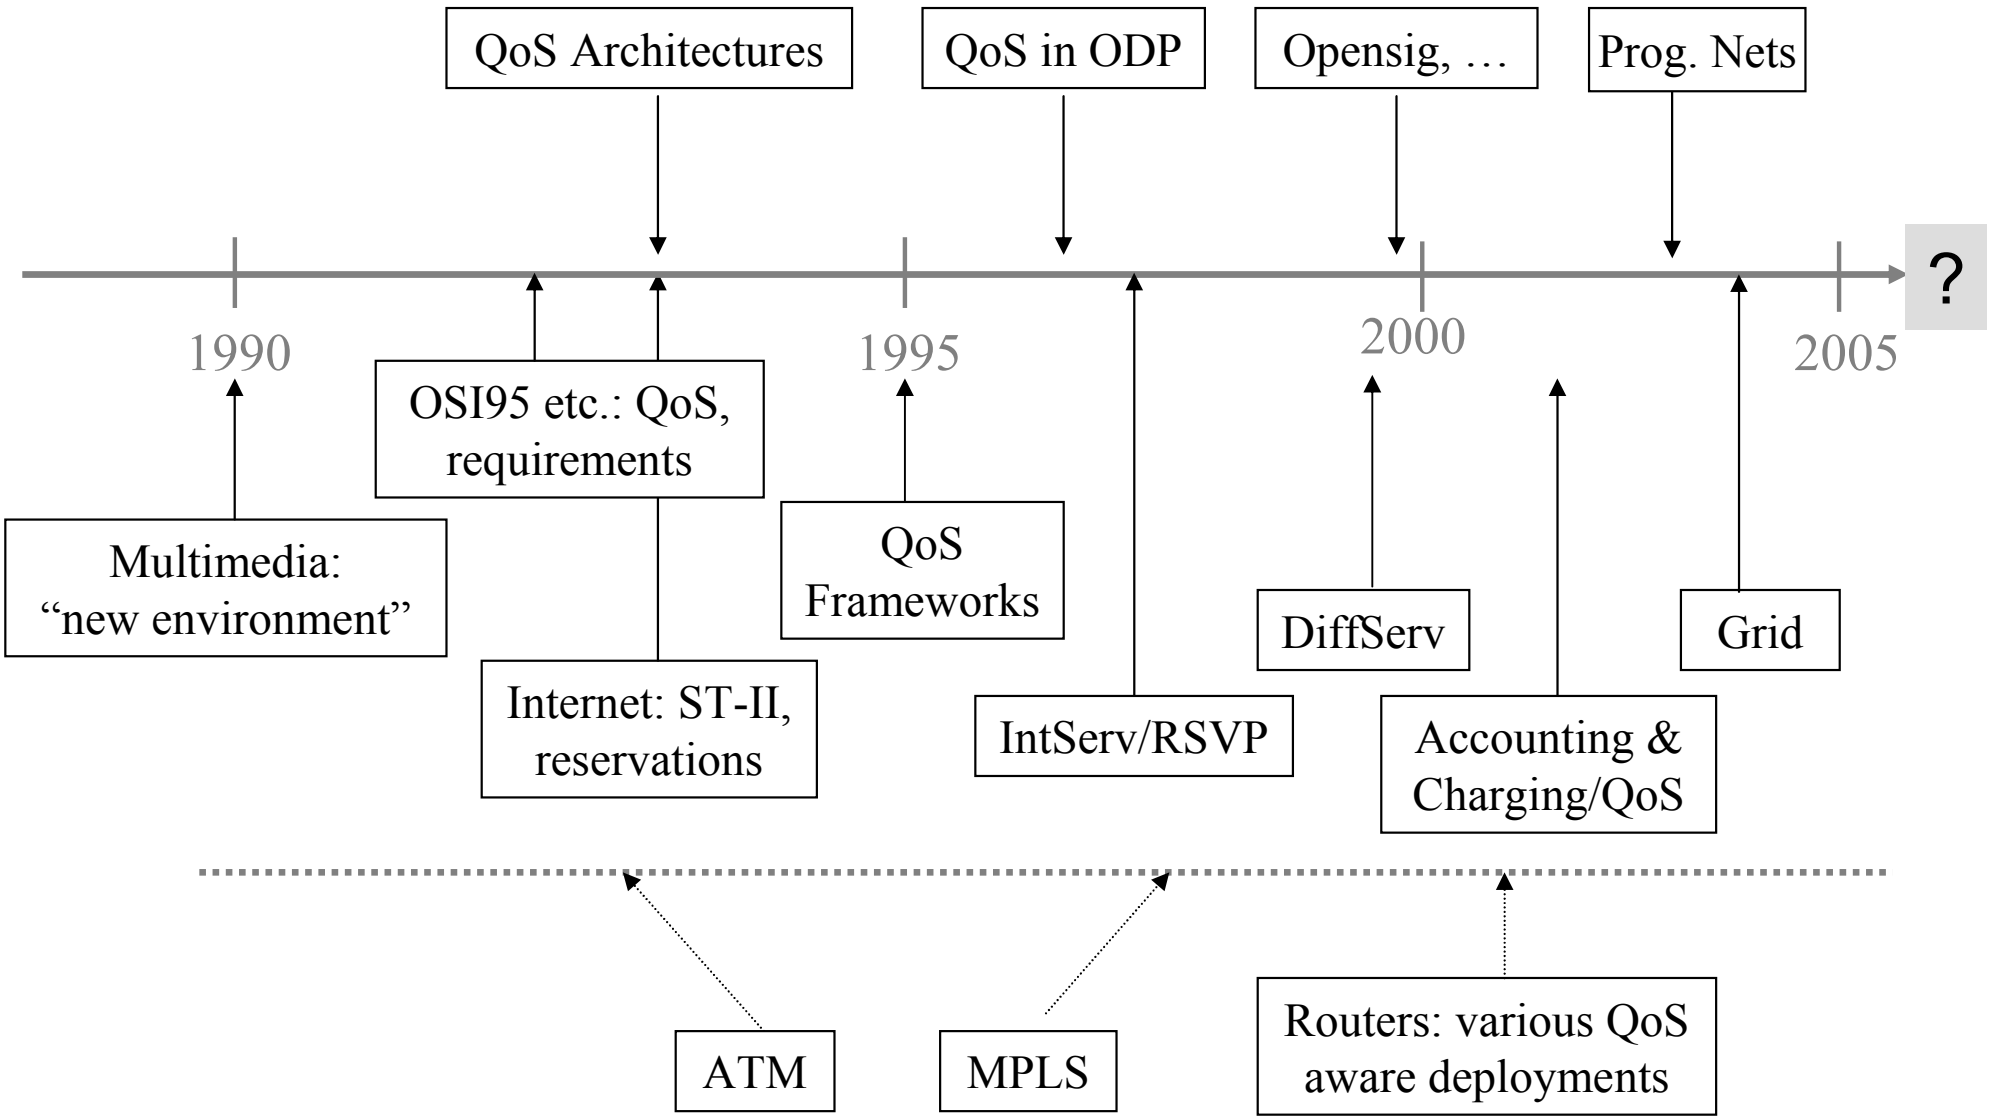

QoS developments, and future?

Developed in the mid 1990s:

Aspects to explore:
performance
availability
security

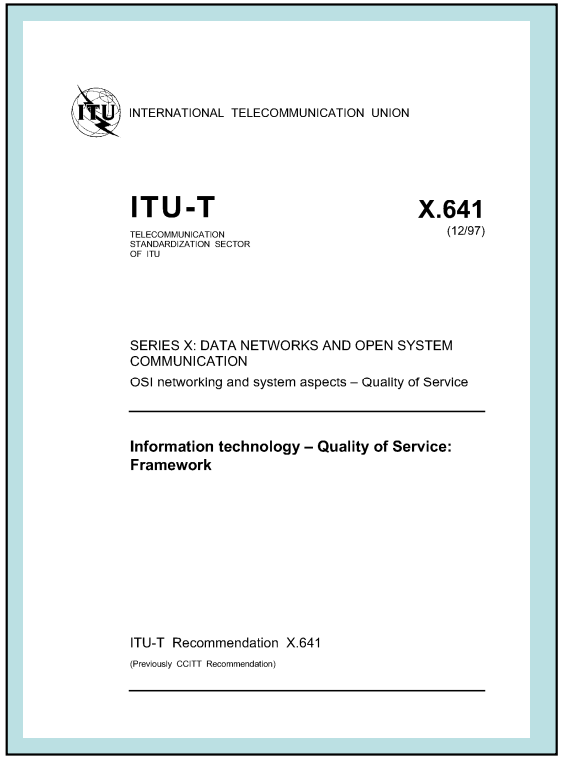

INTERNATIONAL TELECOMMUNICATION UNION

ITU-T **X.641**
(12/97)

TELECOMMUNICATION STANDARDIZATION SECTOR OF ITU

SERIES X: DATA NETWORKS AND OPEN SYSTEM COMMUNICATION

OSI networking and system aspects – Quality of Service

Information technology – Quality of Service: Framework

ITU-T Recommendation X.641

(Previously CCITT Recommendation)

Relationship between QoS aspects:

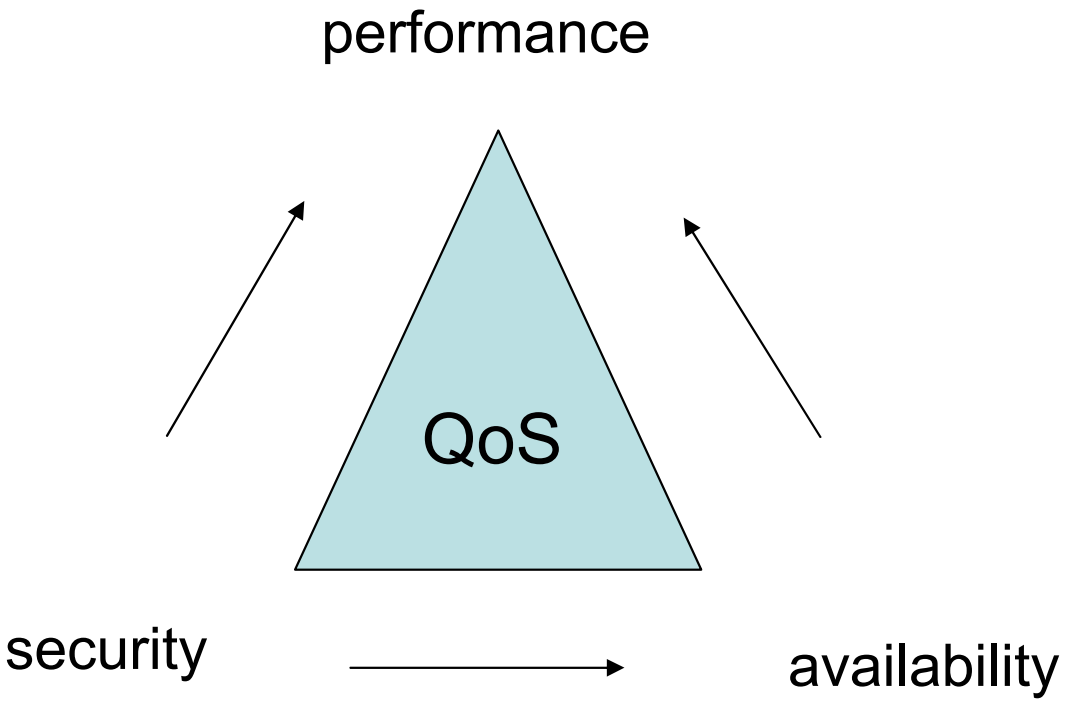

The QoS triangle

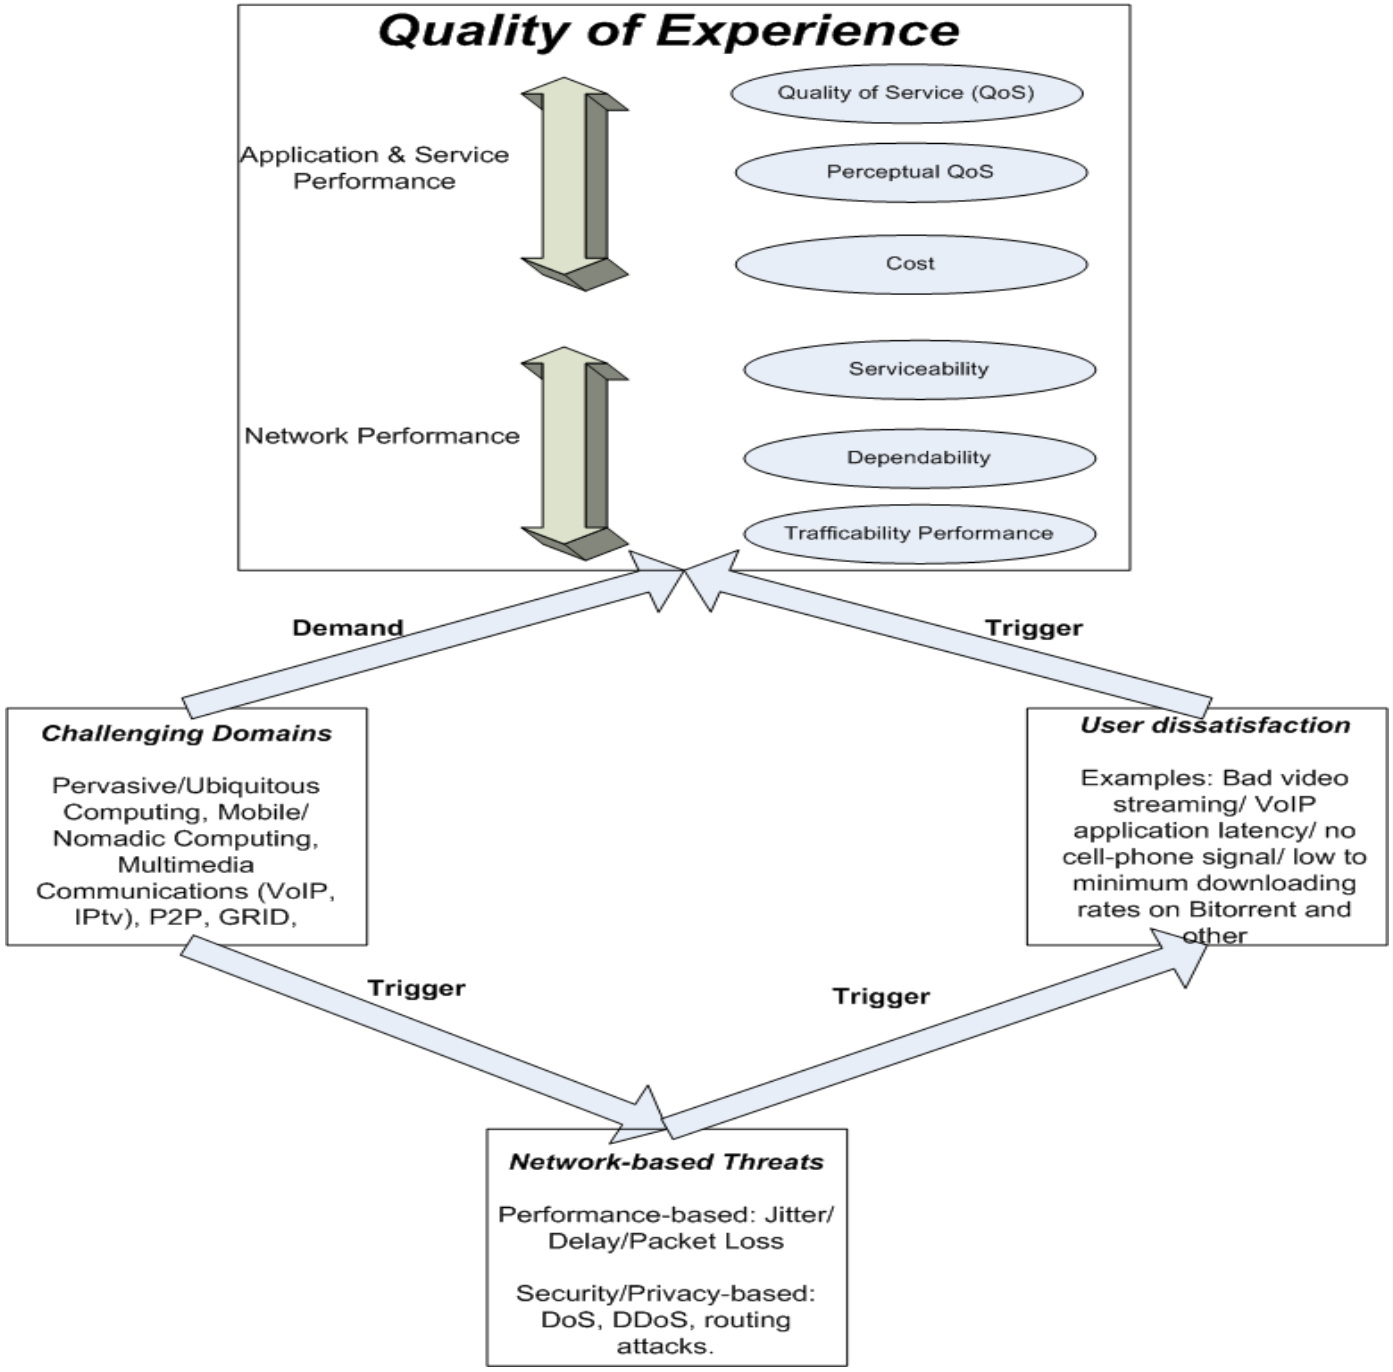

Perceptual QoS, or QoE: the end-user as king

Traffic priority classes agreed during ATM days, for DAV etc.

Why shouldn't network users always ask for highest priority?

Pricing for congestion control?

QoS as a market differentiator

Ref: Internet Economics, MIT Press, May 1997
Edited by Lee W. McKnight and Joseph P. Bailey

Current technical context for QoS:

1. Modern routers' QoS toolset:

   Packet marking, RED, ECN
   DiffServ, IntServ, MPLS
   Inter-AS and BGP …

2. Policy choices / application:

   Dependence on h/w toolset
   and on s/w, eg Cisco IOS

UK-based QoS trials:

1. Core network
2. Access networks

Outcomes:

1. No need to use QoS mechanisms because the network is so heavily over-provisioned
2. Different implementations may have somewhat different capabilities, sometimes heavily limiting the number of QoS classes that can be supported

Ref: http://www.ja.net/development/network-engineering/qos/qos-project-details.html

Recent visits to telcos / ISPs:

Management / control crucial
Care in physical layer set-up
Human always (?) in the loop
Monitoring tools are essential

Traffic is on the increase
Keep the customers happy …
Manage expectations, investment
Dealing with failures, attacks and so on

Current business context for QoS:

QoS-ready routers

Business networks vs ISP

Bottlenecks – do the solutions lie
in the technical or economic domain?

Net Neutrality ... Much debated …

Ref: Net Neutrality: The Technical Side of the Debate
Jon Crowcroft, ACM SIGCOMM CCR, January 2007

Does QoS have a role?

Broadband access future?

100 Mb/s to the home

driven by (two-way) HDTV (x 3) + …

Provide by means of Fibre To The Home/Curb?

What about wireless access, smart phones etc ..?

Recent media articles:

BBC iPlayer

Bandwidth overload?

Cicconi (AT&T): traffic explosion

Traffic interference by ISPs, eg with P2P

Internet overload?

Should QoS disappear as a topic?
Has the subject run its course?

(Should IWQoS disappear ..?)

QoS has a place in other events

That cover, for example:

_ Performance engineering

_ Network/systems management

_ Middleware / distributed computing

But further QoS research remains to be done in areas such as wireless networks and in DTNs where resources are scarce. Also, we surely still need a decent business model!

So, areas of QoS still to be explored:

Performance for mobility / wireless and other 'scarce resource' areas

Resilience for business and for many other critical applications!

And also QoS deployment* …

*See: Geoff Huston (in QoS Fact or Fiction, March 2000):
"More effort is required to turn a QoS Internet into a reliable production platform" – still true today

Further QoS research?

The argument for resilience as a key QoS research area:

_ Attacks are on the increase; traffic is on the increase

_ There is now a considerable dependence on networks and networked systems as a critical infrastructure for the Digital Economy

_ Networks and services **need** to be very highly available

"It should be noted that 2007 turned out to be the most 'viral' year in history. The total number of IT threats more than doubled during the year. In 2007, Kaspersky Lab added almost as many signatures to its databases as it had during the preceding 15 years." - Kaspersky Labs

Resilience is the ability of a system to:

Provide an acceptable level of service despite challenges to normal operation

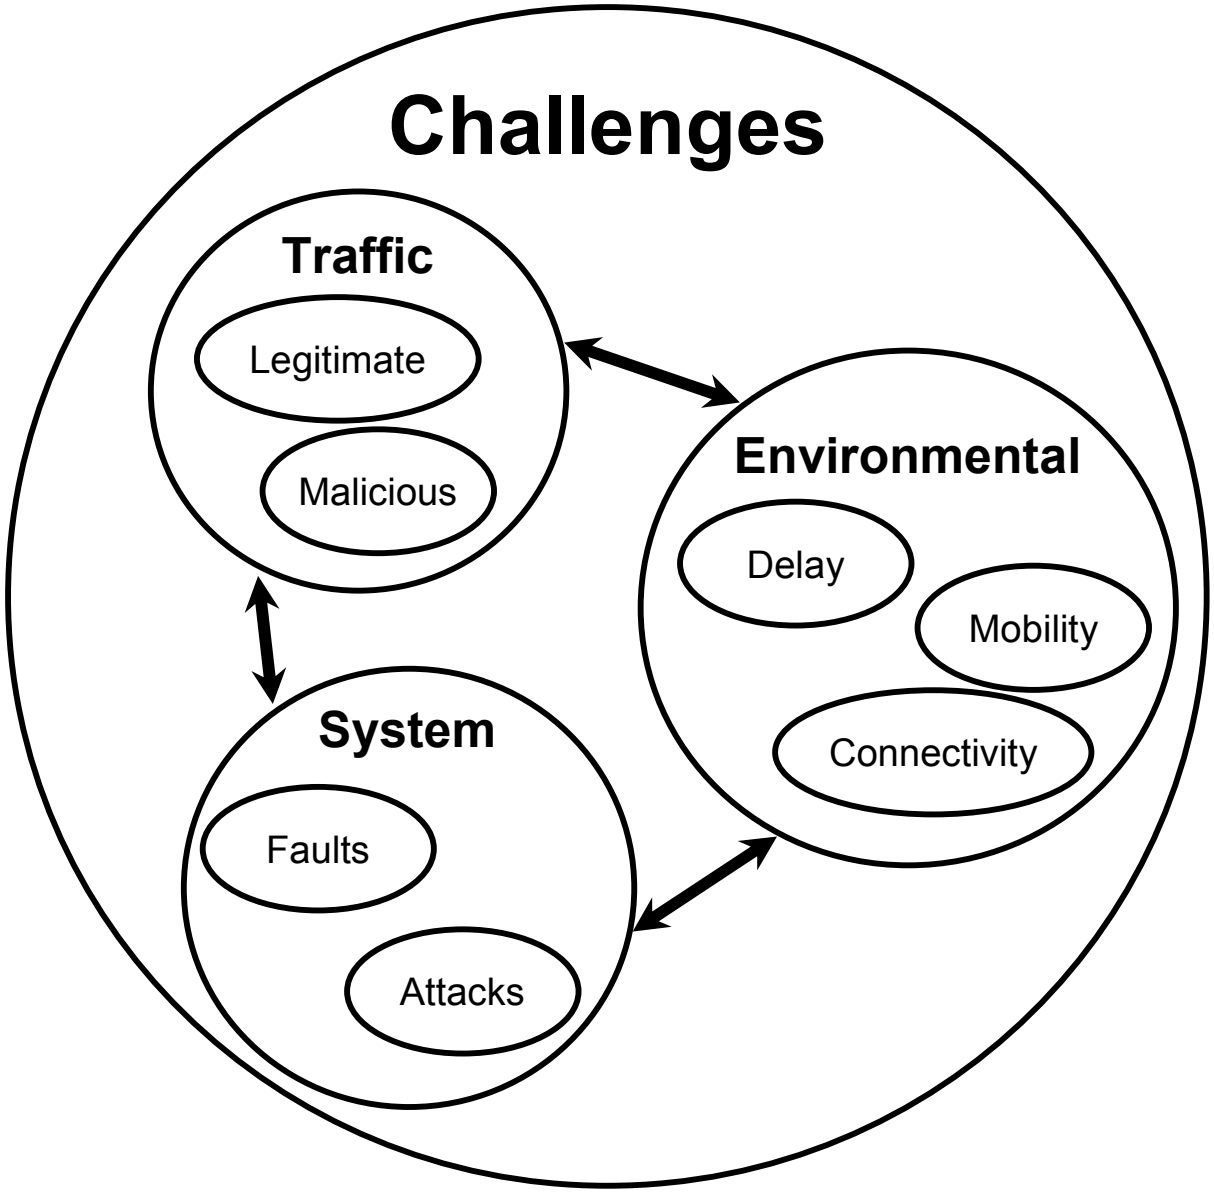

Challenges to normal operation

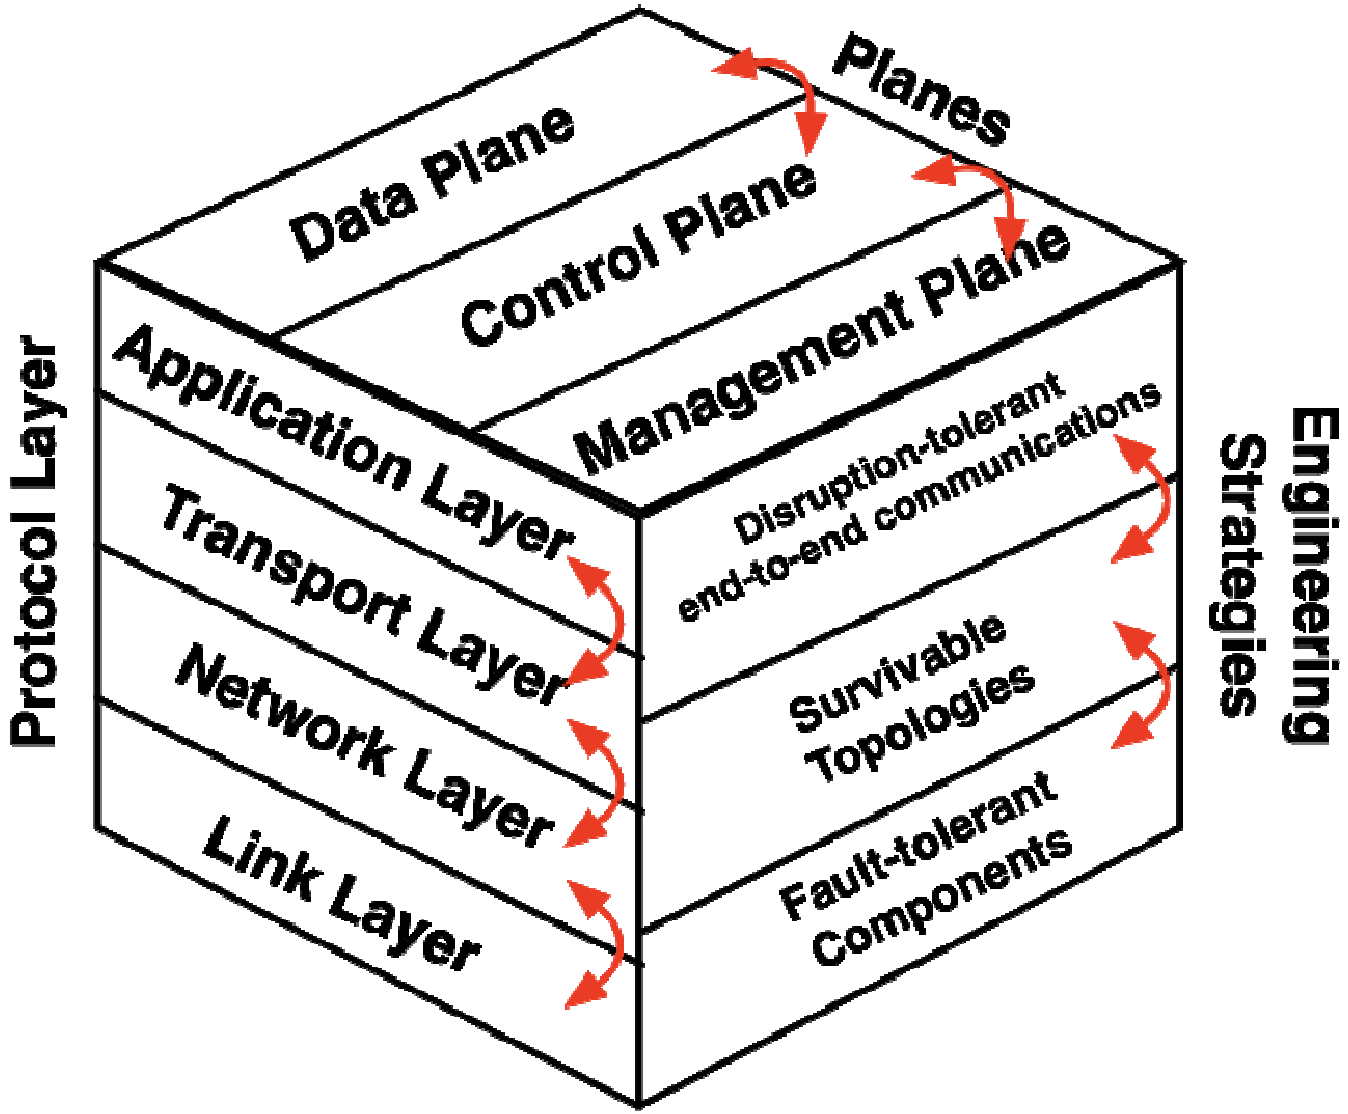

Multi-level resilience

What is normal behaviour

A model for resilience

Identifying anomalies

Choosing metrics

Previous research helps:
especially fault tolerance

The Resilinets project:
Kansas U / Lancaster U
(J Sterbenz & D Hutchison)

Real-time Control Loop

- Defend (proactively)
- Detect
- Remediate (reactively)
- Recover

System Enhancement

- Diagnose
- Refine

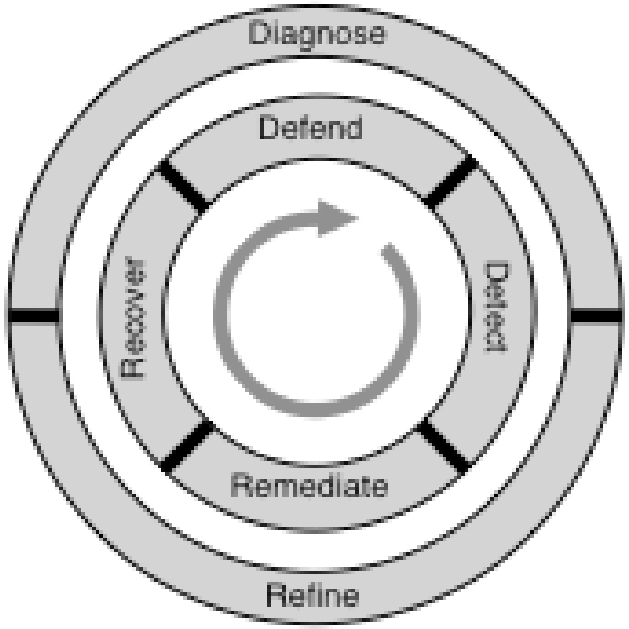

Ref: https://wiki.ittc.ku.edu/resilinets_wiki/index.php/Main_Page

Alternatively:

$D^2R^2 + DR$

could become

$D + DR + R (+ DR)$

0 1 2 3

for phasing of the research

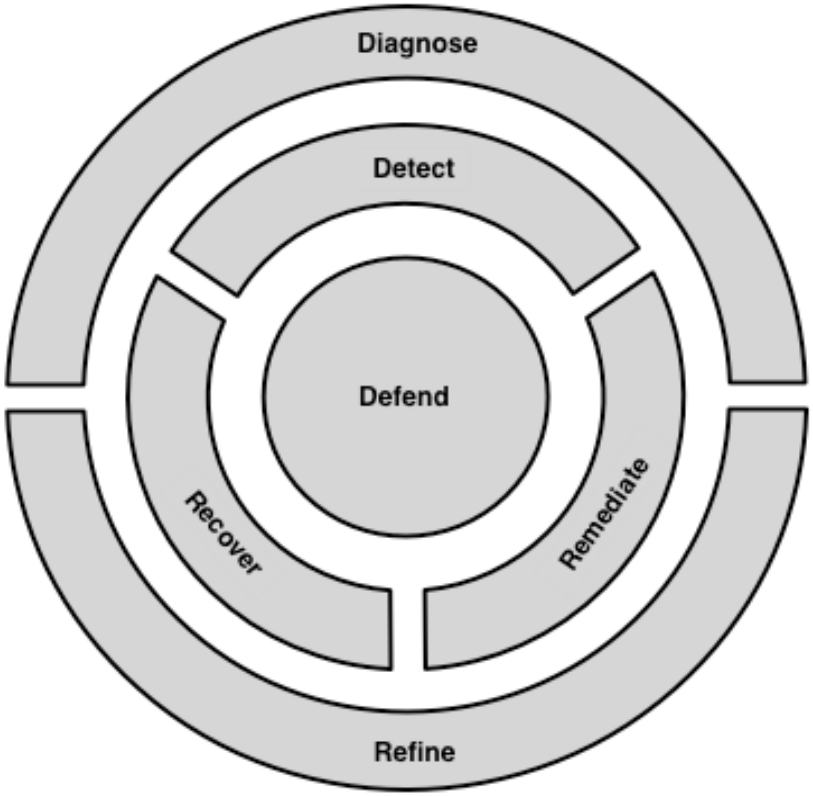

For example, phase 1:

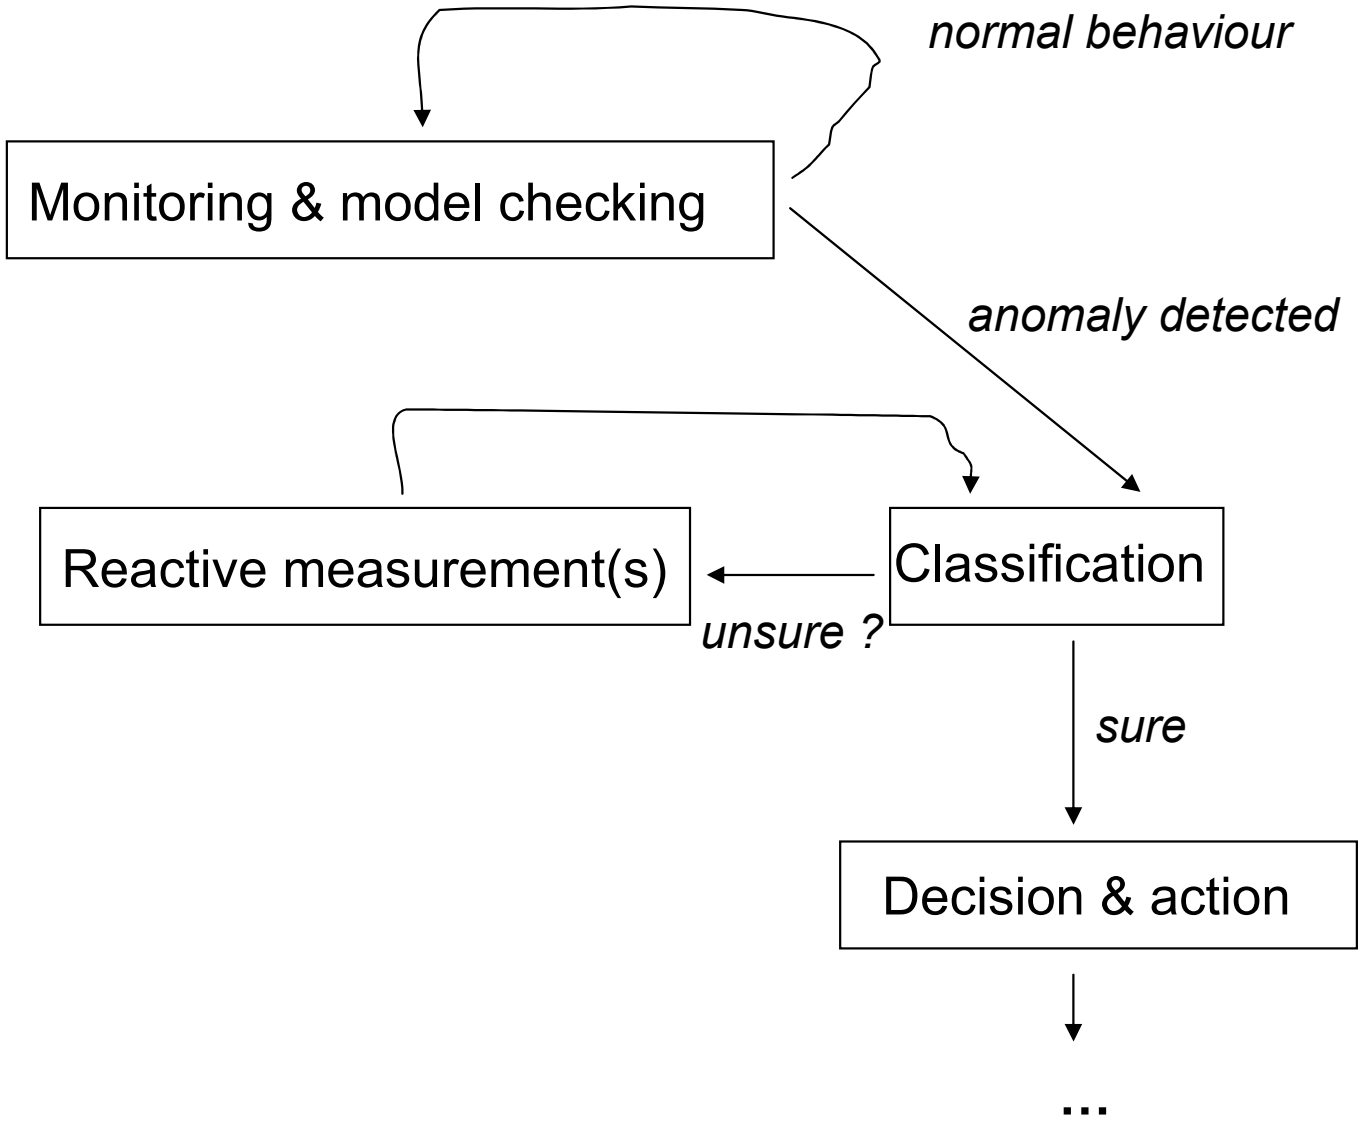

T (the QoS toolset) is what provides aspect P (performance)

R (resilience mechanisms) will contribute to A (availability)

S (the security mechanisms) contribute to both P and A

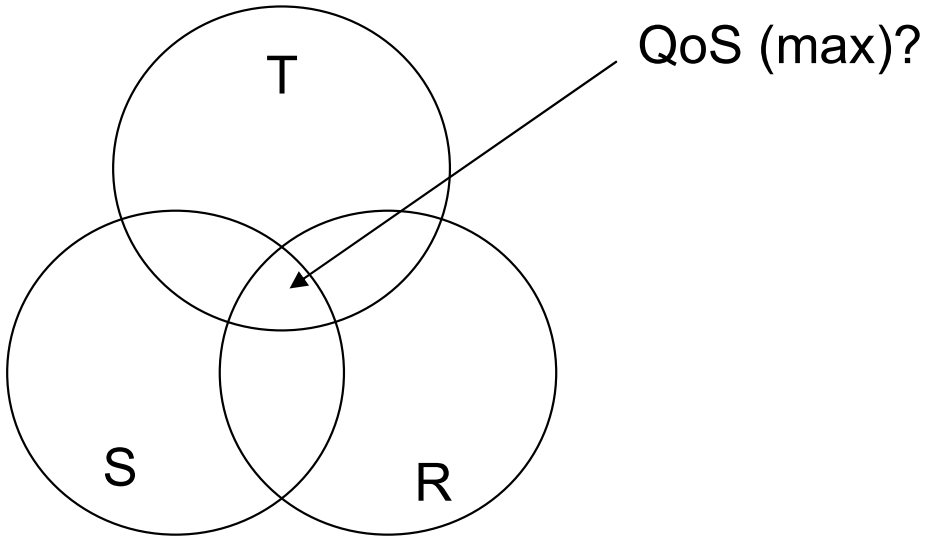

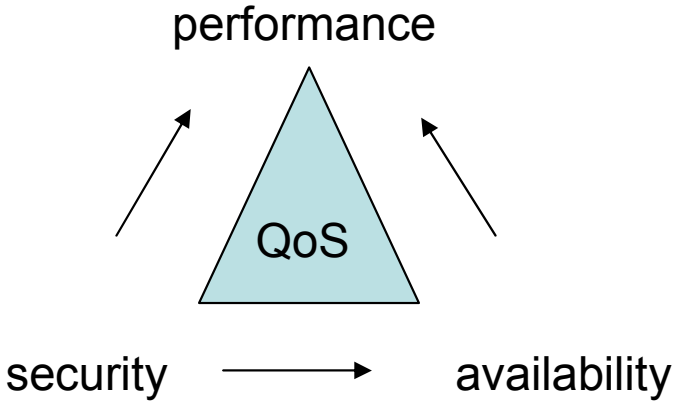

A mechanisms view of the QoS triangle

Availability is an aspect / observable
Resilience is a set of mechanisms?

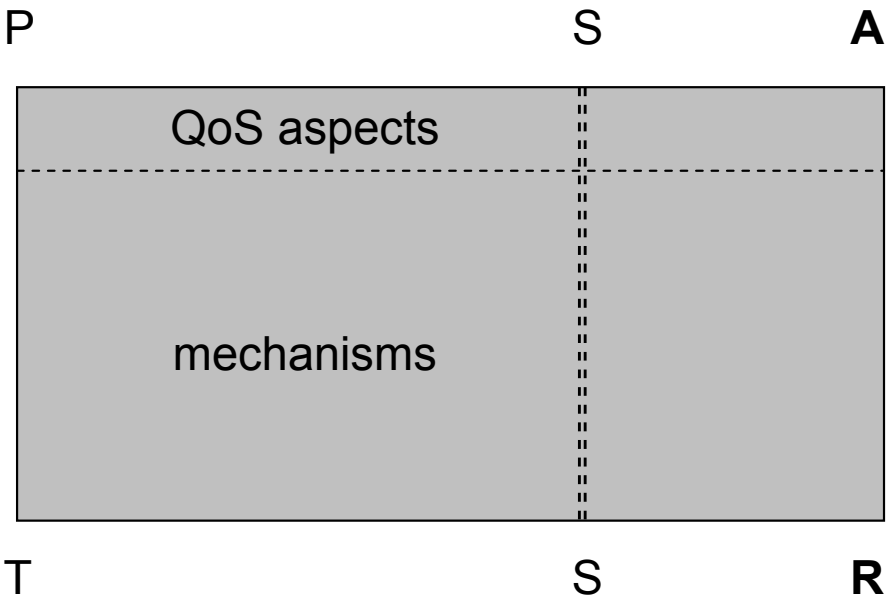

Security as part of the Defend activity?

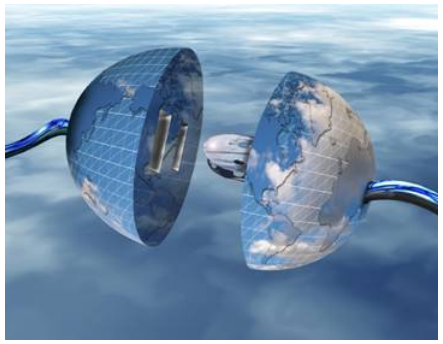

SEVENTH FRAMEWORK PROGRAMME

Resilience and Survivability for future networking: framework, mechanisms, and experimental evaluation

A 3-year FIRE project, to start 1.9.2008

FP7 ResumeNet STREP

LANCASTER UNIVERSITY
InfoLab21

The Internet has become a critical infrastructure – but has it been designed to be one?

The Internet is vulnerable ...

- to flaky communication channels (supporting mobility)
- unintentional misconfiguration
- large scale (natural) disasters
- malicious attacks
- unusual usage and traffic loads

What is needed: A new architectural approach towards a **resilient** Internet

ResumeNet will:

Systematically embed resilience into the future Internet

Three dimensions:

- Conceptual framework
- Mechanisms and algorithms
  - Network resilience
  - Services resilience
- Experimentation in testbeds
  - {network, service, failure, resilience mechanism}

Link with other projects in the Future Internet area

Challenges are inevitable …

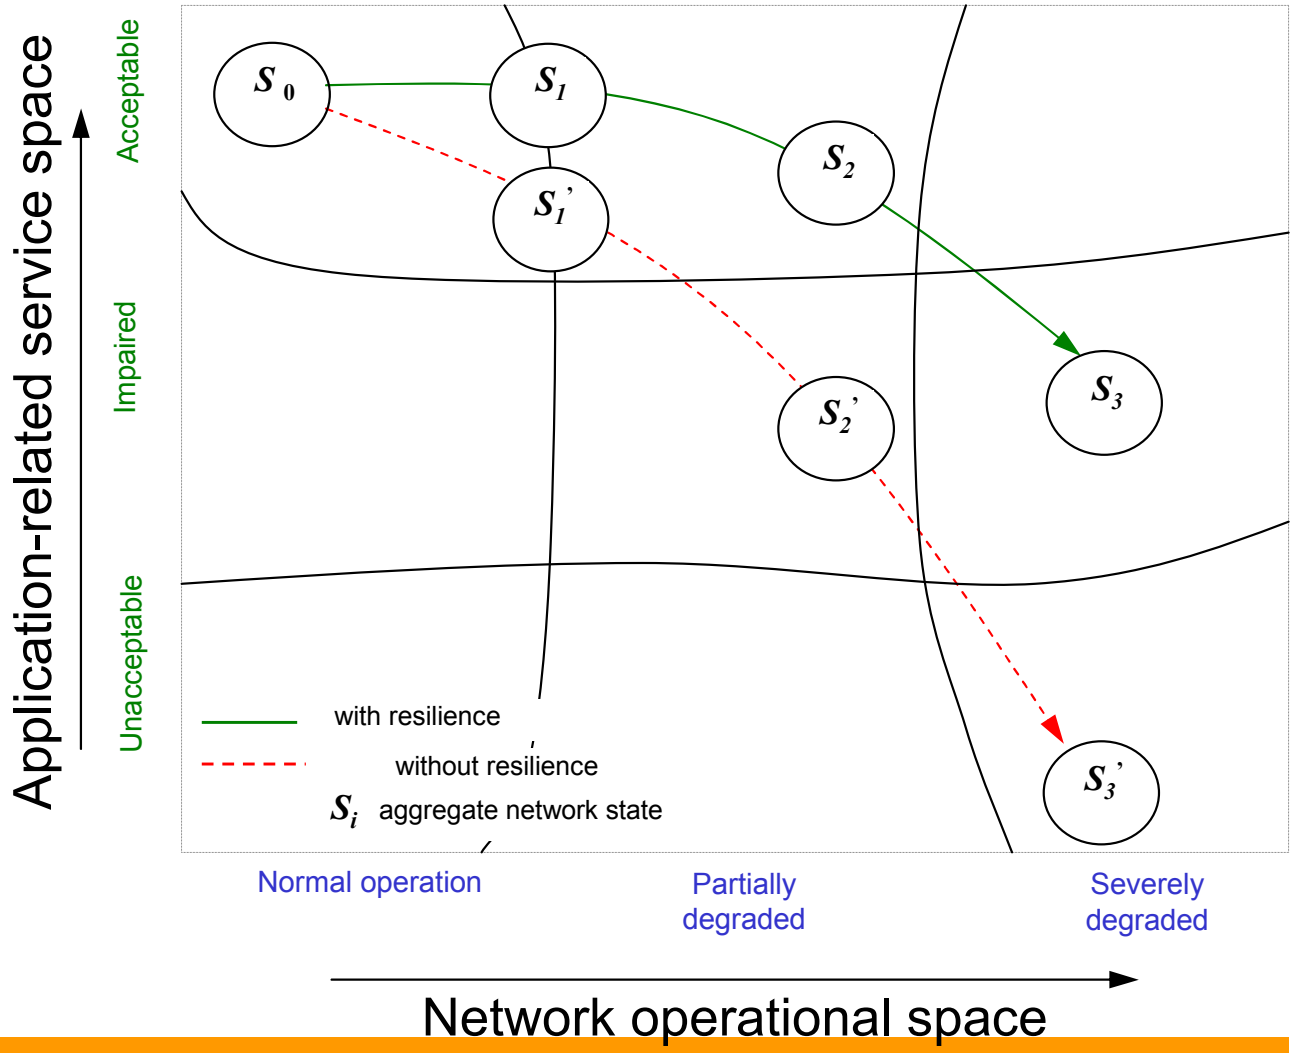

Network and service resilience objectives

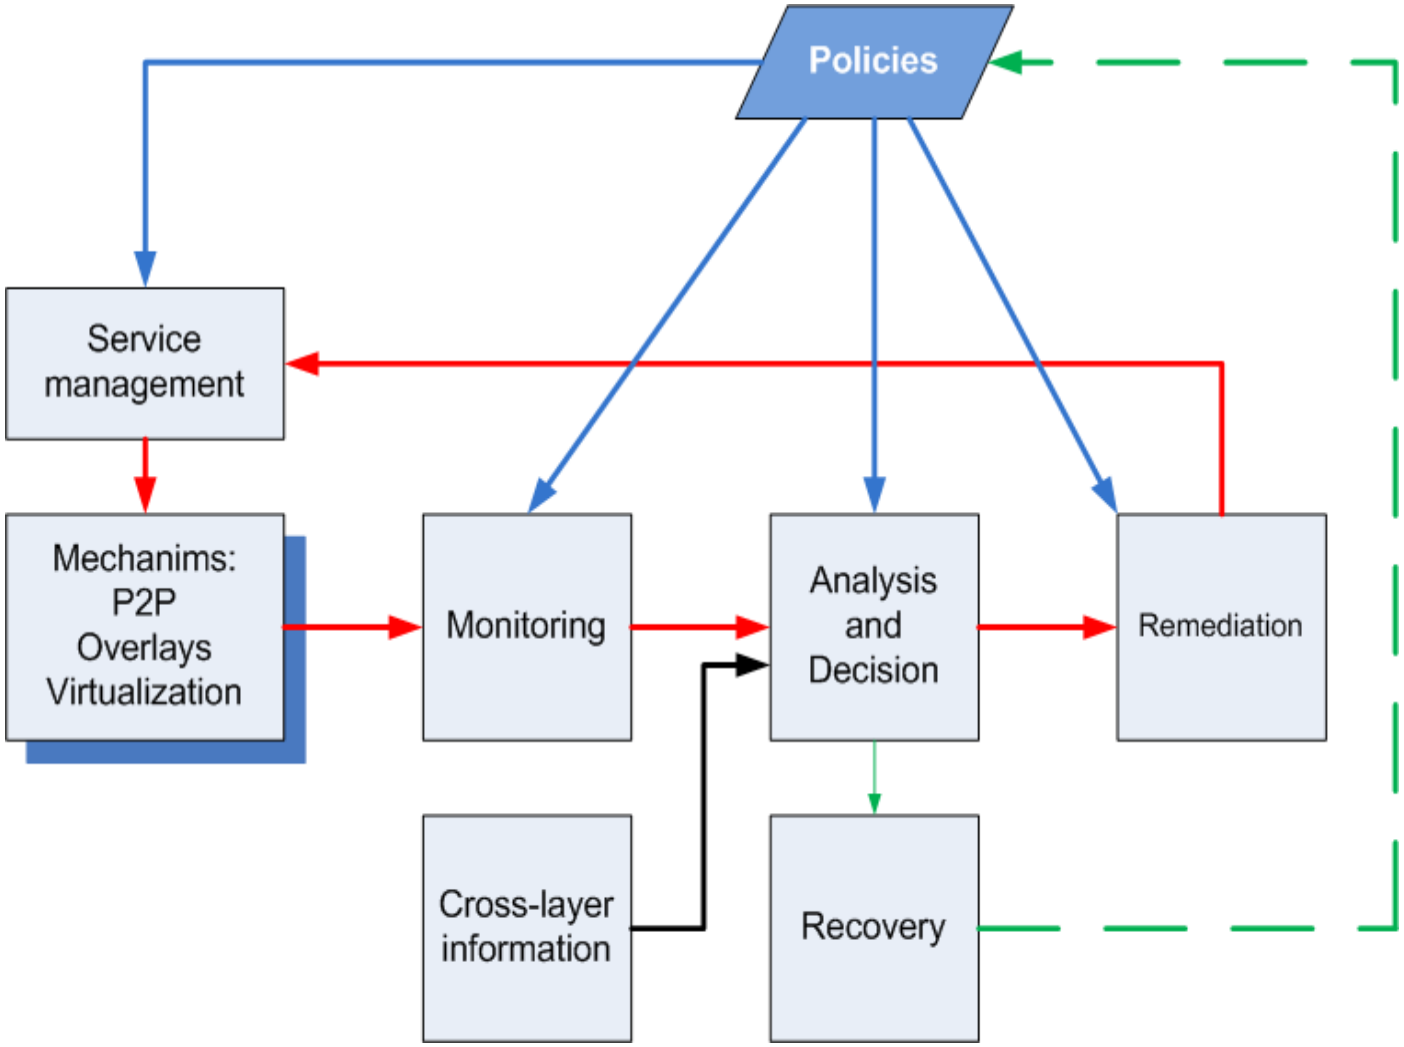

Service resilience architecture

Policies build on resilience metrics work:

Realm Policies

_ Similar in nature to SLAs, but for resilience

_ *Task*: define a policy language

_ *Issue*: resolving conflicting realm policies

Entity Policies

_ How to use resilience capabilities of devices

Mechanism Policies

_ Define applicability context, dependencies, and usage scenarios and outcomes

Service Policies

_ Define desired resilience requirements of services

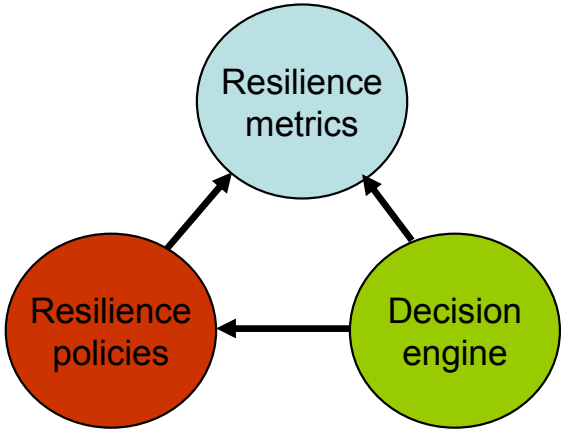

WP structure:

Concepts and framework*
Network resilience
Service resilience
Experimentation / Testbeds
Dissemination
Management

*ResumeNet evaluates & validates D^2R^2+DR strategy and provides guidelines for practitioners

ResumeNet consortium:

ETH Zürich (co-ord)
Lancaster University (*)
Tech. University Münich
France Telecom
NEC Europe Ltd
Universität Passau
Tech. University Delft
Uppsala Universitet
Université de Liège

(*) + U. Kansas and U. Sydney

Project organization

Remember that people are involved:

Security as an abstract concept

Resilience regarded as a ‘tax’

Users pay for performance?

Ref: The Psychology of Security, Ryan West (CACM, April 2008)

To conclude:

QoS is still a considerable challenge:
given today's demanding context
and the increasing dependence
on computer networks

and

_ Performance QoS will still be in demand

_ Resilience needs to be further explored

_ We need to take security more seriously

Thank you! Questions …

From ACM Tech News

Internet Gridlock to Occur in Just Two Years
ZDNet UK (04/21/08) Donoghue, Andrew

Without significant new investment, the Internet's current network architecture will reach the limits of its capacity by 2010, warned AT&T's Jim Cicconi at the Westminster eForum on Web 2.0 in London. "The surge in online content is at the center of the most dramatic changes affecting the Internet today," Cicconi says. "In three years' time, 20 typical households will generate more traffic than the entire Internet today." Cicconi says at least $55 billion in investments are needed in new infrastructure over the next three years in the United States alone, and $130 billion worldwide. The "unprecedented new wave of broadband traffic" will increase fifty-fold by 2015, Cicconi predicts, adding that AT&T will invest $19 billion to maintain its network and upgrade the core of its network. Cicconi adds that more demand for high-definition video will put an increasing strain on the Internet's infrastructure, noting that eight hours of video is loaded onto YouTube every minute, and that HD video consumes seven to 10 times more bandwidth than normal video. "Video will be 80 percent of all traffic by 2010, up from 30 percent today," he says.

http://technology.timesonline.co.uk/tol/news/tech_and_web/article3716781.ece

BBC iPlayer risks overloading the internet

The success of the BBC's iPlayer is putting the internet under severe strain and threatening to bring the network to a halt, internet service providers claimed yesterday. They want the corporation to share the cost of upgrading the network - estimated at £831 million - to cope with the increased workload. Viewers are now watching more than one million BBC programmes online each week. The BBC said yesterday that its iPlayer service, an archive of programmes shown over the previous seven days, was accounting for between 3 and 5 per cent of all internet traffic in Britain, with the first episode of The Apprentice watched more than 100,000 times via a computer. At the same time, the corporation is trying to increase the scope of the service. It is making its iPlayer service available via the Nintendo Wii, allowing owners who are unable to stop playing in time for their favourite programmes to catch up with them later.

Tiscali, the internet service provider, said that the BBC and other broadcasters should "share the costs" of increasing internet capacity to prevent the network coming under strain. The problem for Tiscali, though, is that its concerns are not widely shared in the industry. BT, which provides a key part of the UK's internet infrastructure, said that the problem, "while real", could be solved. It said that the key was not speeding up connections to people's homes, but through improvements in "backhaul and core networks" - the links that operate up and down the country. The iPlayer service has rapidly become a hit after it was introduced at Christmas, even though it involves either watching a programme on a computer screen or finding a way to link the computer to the television. There were 17.2 million requests to watch programmes last month, an increase of 25 per cent on February.

Net Neutrality Battle Returns to the U.S. Senate
CNet (04/22/08) Broache, Anne

At a Senate Commerce Committee hearing on Tuesday titled "The Future of the Internet," Democratic lawmakers argued for a bill that would prohibit broadband operators from creating a "fast lane" for certain types of Internet content and applications. The proposal was criticized by the cable industry, Republican politicians, and FCC Chairman Kevin Martin, who argued that there is no demonstrated need for such action at this point. Much of the discussion at the hearing focused on whether the FCC already has sufficient authority to take action against network operators that interfere unreasonably with their customers' Internet use. Comcast argued that the federal agency does not, while Democrats said their legislation is necessary to clarify the FCC's enforcement role. "To whatever degree people were alleging that this was a solution in search of a problem, it has found its problem," said Sen. John Kerry (D-Mass.). "We have an obligation to try and guarantee that the same freedom and the same creativity that was able to bring us to where we are today continues, going forward." Martin said the FCC does not need to write new regulations because it already has the authority to enforce its existing broadband connectivity principles, which say consumers have the right to access the lawful Internet content and applications of their choice.

[CAnet - news] Glasnost Internet: The threat of Transparency and Privacy to the Internet

From Bill St Arnaud's news summaries

Comcast, Cox Slow BitTorrent Traffic All Day

http://tech.slashdot.org/tech/08/05/15/2028243.shtml

"A study by the Max Planck Institute for Software Systems found that Comcast and Cox Communications are slowing BitTorrent traffic at all times of day, not just peak hours. Comcast was found to be interrupting at least 30% of BitTorrent upload attempts around the clock. At noon, Comcast was interfering with more than 80% of BitTorrent traffic, but it was also slowing more than 60% of BitTorrent traffic at other times, including midnight, 3 a.m. and 8 p.m. Eastern Time in the U.S., the time zone where Comcast is based. Cox was interfering with 100% of the BitTorrent traffic at 1 a.m., 4 a.m. and 5 a.m. Eastern Time. Comcast spokeswoman Sena Fitzmaurice downplayed the results saying, 'P-to-p traffic doesn't necessarily follow normal traffic flows.'"

Evolution of Internet Powers Massive Particle Physics Grid
Network World (04/22/08) Brodkin, Jon

Uncovering clues about the universe's origins is one of the purposes of the Large Hadron Collider (LHC), and distributing the massive volume of data generated by particle collisions to the thousands of researchers around the world is the job of the Worldwide LHC Computing Grid, which will be composed of approximately 20,000 servers. "It's using some advanced features and new technologies within the Internet to distribute the data," says Open Science Grid executive director Ruth Pordes. "It's advancing the technologies, it's advancing the [data transfer] rates, and it's advancing the usability and reliability of the infrastructure." Raw data produced by the collisions is relayed over dedicated 10 Gbps optical-fiber connections to the CERN Computer Center, and from there routed to tape storage as well as to a CPU farm that processes information and generates "event summary data." Eleven Tier-1 sites around the world are then sent subsets of both the raw data and summaries; each site is linked to CERN though a dedicated 10 Gbps connection, while a general purpose research network is used to connect Tier-1 facilities to each other. Once reprocessed by the Tier-1 centers, the raw data is circulated to Tier-2 centers for analysis by physicists via general purpose research networks. Brookhaven National Laboratory's Michael Ernst says the LHC collisions will generate 10 petabytes to 15 petabytes of data annually.

Seven Nato nations have backed a new cyber defence centre in Estonia, which last year blamed Russia for weeks of attacks on its internet structure.

Germany, Slovakia, Latvia, Lithuania, Italy and Spain will staff and fund the hub in the Estonian capital Tallinn.

Estonia came under cyber attack in 2007 after its decision to remove the bronze statue of a Red Army soldier from the centre of Tallinn.

Moscow denied involvement in the flood of data which crashed computers. "We have seen in Estonia that a cyber attack can swiftly become an issue of national security," Nato spokesman James Appathurai said after a signing ceremony in Brussels. "Cyber attacks can cripple societies."

The US will initially send an observer to the project, which will have some 30 staff when fully operational in August.

The centre will provide research, consultation and training on the development of cyber defences for participating national governments.

BBC News, Wednesday, 14 May 2008 17:00 UK

QoS on JANET: Technical Guide
The JANET QoS Development Project

QoS deployments involve a combination of the following components:

- packet scheduling (queuing);
- traffic classification;
- traffic policing and shaping;
- active queue management;
- resource reservation/provisioning and admission control.

Over-Provisioning

Based on the experience of the individual partners during the QoS project, there is an increasingly compelling case to support over-provisioning as the most practical approach to QoS deployment in the Regional Network. While this does not enforce the end-to-end model of QoS provisioning, the natural over-provisioning that is typical in Regional Network cores combined with the potential issues related to vendor equipment makes it the most appropriate near-term solution for supporting QoS. This approach also reinforces the JANET QoS policy which recommends that while no explicit QoS provisioning is adopted, the network should be made QoS-neutral such that traffic marking is not handled separately but is not dropped or altered either. In this case, the Regional Network border routers can also provide admission control based on the packet source/destination address and traffic marking but no further action is taken thereafter.

QoS Traffic Classification

As in most networks, a number of broad traffic types can be identified based on the typical load seen on the Regional Network:

Data transfer - Besides normal web traffic, end sites frequently need to upload and exchange content with servers located elsewhere in order to backup data to an external source or download new content. This is not a critical operation nor is it sensitive to network conditions and thus can be given a low priority.

Web traffic - This class represents normal web requests from machines connected to the network. We can expect this class of traffic to represent the majority of the 'background' load on the system during normal working hours. This traffic is fairly important but can still be classified as normal priority.

VoIP traffic - VoIP is increasingly being used both on a personal level and more formally as a research tool in end sites. Audio traffic typically requires around 64 Kbps and is sensitive to network conditions. As such, VoIP traffic should be given a higher priority than normal traffic where possible.

Video Conferencing traffic - Increasingly, both large and small sites may wish to hold video conferencing calls with other sites regionally, nationally, and beyond. A video conferencing call typically needs around 320 kbps in both directions for video traffic and 64kbps in both directions for the audio traffic. This traffic needs to be highly prioritized due to the strict bandwidth requirements and its delay sensitive nature.

Network Control traffic - Routing information and other network control traffic should be classified as high priority as it is used to exchange state information between routers and other network devices. However, this usually utilizes only a limited amount of bandwidth a so introduces only a minor overhead.

The marking scheme adopted by Lancaster which supplemented this with an additional class specifically to identify Network Control traffic:

Traffic Class	DSCP Value	Binary Value
Network Control	34	100010
Premium	26	011010
Best Effort	18	010010
Less than Best Effort	10	001010

The need to simplify the DSCP marking scheme was further highlighted by issues encountered on some equipment while configuring QoS. We found that several implementations do not offer more than 2-3 separate queues for traffic handling which obviously limits the number of unique traffic classes that can be enforced. As such, our experience is that a limited number of classes makes sense, both in terms of simplifying the QoS policy and ensuring it can be practically enforced on the network.

On Cisco IOS, there are various ways to mark traffic. Marking can be performed using ACLs, such that traffic matching a named ACL can be marked with a given DSCP value, or class-based marking can also be used. The specifics of implementing per hop DSCP handling will be vendor-specific. Here is an example of configuring Premium IP on IOS:

```
class-map EF
 match ip dscp 46
!
policy-map TEST
 class EF
  bandwidth percent 99
!
interface GigabitEthernet0/1
 service-policy output TEST
```

In IOS there is a single command that can be used to police traffic and take a given action based on the observed behaviour, e.g.

```
policy-map TEST
 class premium-aggregate-1
  police 1000000 10000 10000 conform-action transmit exceed-action drop
```

By dropping excess Premium IP traffic rather than remarking to be BE, it ensures that the Premium IP service either works as intended, or fails. To ship packets as BE that are believed to be handled as Premium IP by the source will only cause problems for the application users. It is usually better to reconsider the application usage, or to change the provisioning, than to pass Premium traffic on as BE.

Policy

The policy enforced by a site will be important in determining how QoS will be deployed and supported. The aim of this policy is to act as a management tool to define how QoS resources are allocated within a site in an unambiguous manner and to act as a guideline for long-term deployment and usage. The QoS policy can also be influenced by external entities such as standards body recommendations (e.g. IETF), existing provider policy from the RNO/JANET, and *de facto* best practice. The convention is to specify three DiffServ classes to represent aggregate traffic classed as shown in the table below;

Traffic Class	DSCP	CoS	Applications
PremiumIP	46	5	VoIP, VC, multimedia traffic
Best Effort	0	0	Web traffic, normal file transfers
Less than Best Effort	8	1	Batch operations, large file transfers

The classifications presented here are for illustrative purposes to demonstrate how this aspect of the QoS Policy could be structured. Moreover, certain classes of QoS application (such as control traffic) may not be included in the above list but still need to be represented in some way.