

Resource Allocation, Control, and Accounting
for the Use of Network Resources

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

0. MANAGEMENT SUMMARY

This paper gives reasons for wanting better sharing mechanisms for networks. It concludes that the challenge of sharing network resources (and for example intercontinental link resources) between groups of users is neither well understood, nor well catered for in terms of tools for those responsible for managing the services. The situation is compared with other fields, both inside and outside IT, and examples are cited. Recommendations for further work are made.

The purpose of this RFC is to focus discussion on particular challenges in large service networks in general, and the International IP Internet in particular. No solution discussed in this document is intended as a standard. Rather, it is hoped that a general consensus will emerge as to the appropriate solutions, leading eventually to the adoption of standards.

The structure of the paper is as follows:

1. Findings
2. Conclusions
3. Recommendations

1. FINDINGS

Issues arising from contention in the use of networks are not unusual. Once connectivity and reliability have been addressed to a reasonable level, bandwidth becomes (or appears to become?) the main issue. Usage appears to have a strong tendency to rise to fill the resources available (fully in line with the principles of Parkinson's Law). Line-speed upgrades have an effect, but with no guarantee of permanently alleviating the problem. Line-speeds are increasing as technology improves over time, but the variations on matters like availability and funding are wide, and users remain avaricious.

Often the situation can appear worse than having to survive in a jungle, in the sense that the strong (even if "good") seem to have little advantage over the weak. It may seem that it is the determined person rather than the important work that gets service.

Most people will have experienced poor service on an overloaded network at some time. To help the end-users, it seems on the face of it that one must help the IT Service Manager he relates to. Examples relating to the relationship between the network manager and his customers, IT Service Managers at institutions connecting to his network, include the following:

(a) If the IT Service Manager finds his link to the Network Manager's network overloaded, he may be offered a link upgrade, probably with a cost estimate. He might prefer control mechanisms whereby he can say that department X deserves more resources than department Y, or that interactive terminal use takes preference over file transfers, or that user U is more important than user V.

(b) Where an IT Service Manager is sharing a link, he will commonly get more than his institution's share of the link, and often get very good value-for-money compared to using a dedicated link, but he has no guarantee that his end-users' usage won't get swamped by the use of other (perhaps much larger) partners on the shared link. This could be seen as wishing to have a guaranteed minimum share according to some parameter(s).

(c) On a shared link as under (b), the Network Manager may wish to ensure that usage of the link (which might be a high-performance trunk line on a network or an international link for example) by any one partner is "reasonable" in relation perhaps to his contribution to the costs. In contrast to (b), the Network Manager is wishing to impose a maximum value on some parameter(s). He may be happy if the width of the IT Service Manager's access link is not greater than his share of the shared link (assuming the measure agreed on is "width"), but this will commonly not be the case. To be able to reach agreement, the Network Manager and the IT Service Manager may need options on the choice of parameters, and perhaps a choice on the means of control, as well as being able to negotiate about values.

In circumstances where the Network Manager can exercise such controls over his customers, the IT Service Managers may say with some feeling and perhaps with justification, that if they are going to be controlled can the Network Manager please provide tools whereby they can arrange for the onward sharing of the resource they have, and thence onwards down the hierarchy to the end-users.

(d) It may be Network Manager A has a link that Network Manager B would like to use on occasion, perhaps as back-up on access to a third network. Network Manager A might well wish to be accommodating, perhaps as examples because of financial benefit or perhaps because of the possibility of a reciprocal arrangement. However, the fear of overload affecting normal use and the lack of control over the usage militates against arrangements that the parties could be quite keen to make.

Such challenges are very far from being unique to networking. Government and both public and private organisations and companies allocate budgets (and resources other than money), control and account for usage, recognising the possibility of overdrawn and borrowing. In times of shortage, food is rationed. I haven't checked this out, but it would surprise me if Jerry Hall wasn't guaranteed a ticket for any Rolling Stones concert, should she wish to attend.

The charging factor influences use but does not control it (except perhaps in unusual circumstances where say payment was expected in advance and usage was cut off when the money ran out).

In the IT world, multi-user hosts have filestore control systems; one that I use has an overdraft facility with no penalty for not having a prior arrangement! There are also system designs and implementations for sharing host processor time with more sophistication than just counting seconds and chopping people off; this problem seems to me to be reasonably well understood. (Library catalogue searches under author "John Larmouth" should provide some references for those who require convincing.) Some multi-user hosts have controls of sorts on terminal connections. On the other hand, I am not aware of any control system in operation that can guarantee multi-user host response time even outside the network context among directly connected terminals.

The various roles bring different interests to bear. A provider will not necessarily see it in his interests to control usage, or (perhaps even more likely) to provide customers with control tools, since the lack of these may encourage - or even oblige - the customer to buy more. Even if the IT Service Manager can deal with the issue of who or what is important, and the issues of the relative importance of allocating resources against requests, other issues like social acceptability may arise to complicate his life. For example it may be generally agreed (and perhaps the network manager instructed) that "everyone" must be able to do a small amount of work at any time, perhaps to do some housekeeping or seek information.

Time is an important factor. Network resources, like computer processor time and unlike filestore, vanish if they are not used. People will in general prefer resources during prime shift to those in the middle of their night; however, in global terms the middle of their night can be during prime shift somewhere along their path of usage.

What's to do? Splitting lines with multiplexers is rather inflexible, and may well militate against the benefits of resource-sharing that give rise commonly to link-sharing arrangements. Some technologies:

- have the ability to treat (or at least mark) traffic as of high priority, for example where it gives emergency or status information;
- (in the case of X.25(84), I understand from my JNT colleague Ian Smith,) have throughput class (section 6.13) and transit delay (section 6.27). (Ian tells me that it is in his view far from clear how practical these facilities are);
- may be able to discriminate between traffic on grounds of network source address;
- may be able to discriminate between traffic on grounds of network destination address;
- may be able to discriminate between traffic on grounds of application protocol, perhaps giving preference to interactive terminal traffic, or making a choice between preference for email and for file transfer traffic;
- may be able to discriminate between traffic on grounds of other facets of network protocol or traffic.

In practice, one may well not have adequate tools in these or other terms, and one may well have to ignore the challenges of resource control, and either ignore the issue or refuse service.

2. CONCLUSIONS

2.1 There seems to be a lack of tools to enable the controlling and the sharing of networks and links. This is militating against the cooperative sharing of resources, and restricting the ability of organisations to do business with one another.

2.2 Further, the definition of what constitutes a share, or what parameter of service one would try to measure and control (or what

the choices are if any), is not clear.

2.3 Following from that, it is then not clear whether what is needed is new or enhanced protocols/services, new or enhanced procurement specifications or profiles, or new or enhanced networking products or tools.

2.4 Service providers (more likely the public carriers or but also some Network Managers) may see it as against their interests to provide controlling tools if they see them as tending to constrain usage and hence reducing income. If so, they may not support, and may even oppose, progress in the area. However, they might be persuaded that the provision of such tools might give them competitive edge over their rivals, and therefore to support appropriate projects and developments.

3. RECOMMENDATIONS

There seems scope for one or more studies to:

- restate and refine the definition of the problems;
- collect, catalogue and relate relevant experience in both the networking and non-networking fields;
- make recommendations as to what areas (e.g., among those suggested in 2.3 above) projects should be undertaken;
- outline possible projects, indicating the timescale on which improved sharing of production network service resources is likely to be achieved, and recommending an order of priority among the suggested projects.

FOOTNOTES:

Gender issues - where appropriate, the male embraces the female and vice versa.

Dramatis Personae:

Jerry Hall is a close associate of Mr. M. Jagger, formerly of the London School of Economics in the University of London, and now Chairman and Chief Executive of an internationally prominent and successful commercial musical operation.

Others mentioned in this paper are assumed to prefer to remain anonymous, although the standard is to give contact information for the author (see Author's Address section).

Security Considerations

Security issues are not discussed in this memo.

Author's Address

Phil Jones
JNT
RAL, Chilton, Didcot, OXON OX11 0QX

Voice: +44-235-446618
Fax: +44-235-446251

Email: p.jones@jnt.ac.uk or c=gb;a= ;p=uk.ac;o=jnt;i=p;s=jones;