Workshop Announcement

Title: Automated Resource Sharing on the ARPANET

Date: Monday May 21, 1973

Time: 9:00 AM to 5:00 PM

Place: Bolt Beranek and Newman Inc., Cambridge, Mass.

Hosts: TENEX and TIP Groups at BBN

Theme:

This workshop will focus on various aspects of the question:

    What steps can be taken to automate access to the distributed
    resources on the ARPANET?

In particular, how can we move from where we are today toward an
environment which facilitates resource sharing by moving the burden of
dealing with the network from the human user to processes which act on
his behalf? Additionally, operating systems themselves perform various
operations not directly initiated by human users which could better be
performed with the availability of resources on other systems (e.g.
file system backup); how can we move toward an environment which
facilitates such system-system cooperation?

Objectives of Workshop:
----

1. To identify and clarify the issues raised by automated resource
    sharing.
    What are the obstacles preventing more widespread resource sharing
    on the ARPANET? Are they technical, political, administrative in
    nature? Is it that there are few resources worth sharing (we don't
    think so)? Is automated sharing a bad idea (We don't think so)?

2. To identify resources at various network sites appropriate for
    automated sharing; and to identify the need for resources which
    don't but should exist.

3. To formulate a series of experiments for the purpose of evaluating
    relative merits and disadvantages of different approaches to
    automating resource sharing.
    The intent of such experimentation is to gain experience through
    construction and use of prototype systems which support automated
    sharing.

Format of Workshop:

Morning:

In order to get the workshop "up to speed", each participant will be
expected to give a brief presentation of relevant work he (his site) is
currently engaged in, is planning to do, or to identify and discuss
issues he feels are relevant to the subject. Time will be allowed for
brief discussion after each presentation.

Afternoon:

General discussion of the issues raised during the morning session.
Possible subjects for discussion include (but need not be limited to):

1. Identification of possible multi-site "services".
    Intersite mail, terminal linking, status information are some
    examples - what are others?

2. Identification of resources appropriate for remote utilization.
    File systems, compilers, on-line query systems, manuscript
    preparation systems are some examples - what are others?

3. Access to remote resources.
    Possibility of access paths other than the standard logger port. To
    what extent (if at all) can the access paths to a variety of
    different resources be standardized? How can resources which may
    move from Host to Host or may be available on several Hosts be
    dynamically located and selected for use? The need for
    (desirability of) a "broadcast ICP".

4. Problems of accounting for resource utilization.
    Some form of network wide accounting would be a great convenience.
    For example, it would be nice if a user could use the same account
    at many (all?) sites. What are the problems (if any) preventing
    this?

5. Problems of security and access control.
    Authentication of users/processes attempting to use resources. As
    with network wide accounts, the ability to use the same name and
    password at all sites would be convenient. How can a user's
    password and other sensitive data be protected in such an
    environment?
    The notion of a third party password validation and user
    authentication service.

6. Approaches to automating resource sharing.
    It is possible without difficulty to identify several which on the
    surface appear to be different:

    a. Multi-site executive programs which make resources accessible to
        the user at the command language level; e.g. the inter-site,
        user-user interaction and file maintenance activity supported by
        the RSEXEC.
    b. A programming language environment designed to facilitate
        resource sharing; e.g. LISP is a machine independent language -
        one could imagine a multi-computer LISP system which supported
        automated resource sharing.
    c. The "collect a resource" approach - identify an Editor here,
        file storage service there, a compiler somewhere else, etc; and
        build a "workshop" environment which provides convenient access
        to these resources.

    What are the relative merits and disadvantages of these approaches?
    What aspects do these approaches have in common? Is it possible to
    identify a common base capable of supporting them all?

7. Protocols to support automated resource sharing.
    It would be inappropriate to attempt to generate a detailed protocol
    specification at this workshop. However, it is appropriate to
    discuss the kinds of activity a protocol should support. Existing
    protocols (excepting Host-Host protocol and possibly, the new TELNET
    protocol) appear to be oriented toward human users. Automated
    resource sharing suggests processes acting on behalf of human users
    to interface to remote resources; this in turn suggests that the
    protocols should be highly process oriented. For example, because
    there should be minimal human intervention in error recovery, the
    protocols should be extremely robust; e.g., include well specified
    time outs, etc.

Arrangements:

If you are planning to attend the workshop, please notify Bob Thomas at
BBN (send net mail to BTHOMAS@BBN, telephone (617) 491-1850, x483). If
you would like us to make motel reservations for you (at the homestead
Inn at Fresh Pond) call Mrs Terry Bernier at BBN (x545).

It is possible that a single day will prove to be insufficient for this
workshop. If that is the consensus of the attendees, the workshop will
continue through Tuesday May 22.

Position papers, memos, notes, etc. prepared by participants in advance
of the workshop will help contribute to the success of the workshop and
are requested. All such papers received before May 11 will be
distributed, in advance, to workshop attendees.

The following questions may be helpful in focusing your thinking:

- What resources would your site be willing to make available for use in
  automated resource sharing experiments?
- Under what conditions would your site be willing or able to
  participate in such experiments?
- What administrative and/or technical considerations would prevent your
  site from entering into a network wide resource sharing agreement?
- If you employ accounting Procedures that require cost recovery, how,
  if at all, should they be modified to work in a network resource
  sharing environment?

Reading List:

We are aware of little that has been written on the subject of automated
resource sharing. However, the following items are relevant (at least
marginally) to the workshop. Please inform us of others of which you
are aware.

1. ARPANET NEWS, Issue 2, Report on COMPCON 73 "Birds of a Feather
    Session" on Resource Sharing Networks, NIC 15337.
2. "A Resource Sharing Executive for the ARPANET", R. Thomas, Preprint
    of paper for 1973 National Computer Conference, BBN Report 2522, NIC
    #14689.
3. "Terminal Access to the ARPANET - Experience and Improvements", N.
    Mimno, B. Cosell, Walden, et. al., COMPCON 73 Proceedings, NIC
    14791.
4. "A Tentative Proposal for a Modified User Protocol", M. Padlipsky,
    RFC 451, NIC #14135.

5. "Interentity Communication - An experiment", R. Bressler, R. Thomas,
    RFC 441, NIC 13773.
6. "Netbank", J. Postel, RFC 408, NIC #12390.

            [This RFC was put into machine readable form for entry]
            [into the online RFC archives by Alex McKenzie with   ]
            [support from GTE, formerly BBN Corp.   9/99]