Expanding International Email Connectivity--Another Look

Expanding International Email Connectivity--Another Look

John C. Klensin and Randy Bush

(Originally published in ConneXions June 30, 1996)

Unlike the situation of thirty years ago, when almost all important scientific work occurred in Western countries, the scientific community is becoming increasingly international. Important work and areas of study occur all over the world. Collaborations and ability to access sources of data and other resources are increasingly important to scientific progress. In many fields, we see more and more inter-institutional collaborations on research and papers that draw on the strengths of each of these institutions. Exchanges of ideas and collaboration and review of proposals should not be limited to one country, or even to developed areas. Especially in such areas as the

```
health
and social sciences and in all of the various fields that study "global
and
environmental future" issues, participation of scientists in developing
areas
has become crucial.  This is true whether the scientists themselves are
indigenous to, or visiting in, those areas; indeed, as the community
becomes
more international, the distinction between the two is gradually becoming
less
clear.

Local policies often reinforce the trend toward collaborations that
require
strong communications links.   For example, as it becomes harder to move
biological samples or historical artifacts across international
boundaries, it
makes increasing sense to do analyses and evaluations within the country
of
origin, then make the data available to both both domestic and remote
parties.
To the degree that relationships involving local and remote scientists
and
institutions become permanent and stable, the benefits tend to spread
with
improved research and more open and tolerant relationships with
regulators and
those not initially involved.

Communication facilities based on computer networks, especially the low-
level
ones such as electronic mail, have become critical to these types of
collaboration.  The post is simply too slow to permit real interaction,
and
fax, while faster, does not lend itself well to group interaction, much
less
true collaboration on shared materials.  The more the network connection
infrastructure can be opened up, the more scientists can participate in
international efforts on the basis of interests, skills, and knowledge,
rather
```

```
than based on where they happen to live or work.

Similarly, data gathered or prepared in remote locations often must be
transferred elsewhere for evaluation or analysis.  The examples usually
cited
are climate or rainfall data, but similar issues arise with health and
nutrition status, mortality and disease figures, and even certain types
of
economic statistics.  Without computer networks, the options are to mail
machine-readable media (often unacceptably slow and unreliable) to send
the
data via fax (typically requiring rekeying or the use of OCR techniques,
which
are not completely reliable), or long-distance international data calls
at
great expense.

Many Latin American, Caribbean, and African universities tell of invited
US
scientists who would not come for sabbaticals or extended research visits
because they would not have email access to their colleagues.  The
problem of
being without network facilities is not merely one of keeping up
correspondence, but of being isolated to the point of ineffectiveness if
reliable email is not available.

On the other hand, when networks are available, previously-unanticipated
collaboration seems to come into being almost spontaneously.  Again, the
underlying causes seem to involve a latent demand that remains latent as
long
as joint work requires either the disruption of waiting for the post, the
continual retyping of texts transmitted by post or fax, or the need to
secure
large budgets and approvals for extensive international travel.

When computer networks are available, people quickly become comfortable
with
using them, the collaboration seems to happen often and quickly in many
disciplines and work groups, and needs only a little bit of outside
stimulus
```

```
to get it started in others.

These patterns in the scientific community have been paralleled by
activities
in various agencies and organizations concerned with development or
assistance.  Groups in one country need to communicate with those in
another,
and, when they can be made available, electronic mail and computer-
assisted
communications have significant advantages over other approaches.

The demand from both communities has been present for some years as made
evident by the "how do I access a network to stay in touch with
colleagues at
home while I'm in the field or to work with colleagues in remote
locations
from my home institution," inquiries which appear at very frequent
intervals
on popular network mailing lists and news groups.  In what we might think
of
as first-generation low-end wide area networking, the response was "call
home": connections from a terminal in the field to a centralized
computer.
The calls might be made with modems over international telephone
connections
or remote-connection PPSDN links (usually X.25), but the essential
communication pattern was remote login to a centralized computer that
hosted
what was, in reality, a centralized email system.  Gradually that type of
arrangement became somewhat less decentralized: a single central computer
that
everyone dialed into was replaced by regional central computers with the
same
type of arrangements and some way of communicating among themselves.

As computer costs dropped and modem technologies improved to permit data
communications over low-quality lines at higher speeds and plausible
costs,
opportunities arose for true computer-computer connections, with people
receiving and composing mail on their local systems, rather than trying
```

```
to
type while connected to remote locations.  Of course, those opportunities
have
been taken advantage of, but often in a way that may inhibit positive
long-term network developments.

Organizations with a need to communicate with subsidiaries or
collaborators
can now establish single-purpose polled or dialout arrangements,
typically
using FidoNet or UUCP technology, that link the components of that
organization, or the collaborators in a particular project, with each
other.
In any given situation, this may be reasonable and the arrangements can
be
established with a minimum of fuss.  Unfortunately, there are also
negative
effects:

 o Participants in one activity tend to become isolated from non-
participants
   and participants in other activities.

 o While inter-country communication may be facilitated, intra-country
   communication may be frustrated: either made impossible altogether or
   forced through very remote gateways.

 o While two separate private arrangements may be cost-effective, the
third
   one rarely is and there are usually major advantages in not starting
the
   second.  If a single organization can afford one polled international
mail
   exchange a day, better service for everyone can typically be arranged
for
   the same costs by sharing resources and arranging multiple exchanges.
If
   the user base is wide enough and can be expanded without regard to
project
   boundaries, it has been shown time and time again that a minimal
```

```
networking
   channel will quickly build a user base which soon fills that channel.
Even
   in less-developed areas, the perception of value rapidly builds to the
   point that the user base then manages to find its own funding to
   continually increase the available bandwidth.

 o Private arrangements tend to satisfy, and then hide, demand.
Resources
   being invested that could contribute to higher-quality connections for
an
   area do not appear when needs or market surveys are performed.  This
is
   especially critical in situations in which the precondition for
building a
   network infrastructure is the ability to demonstrate that the demand
and
   users exist.  This demand often does not arise until networks are
actually
   seen to be in place and working, and people involved in private
network
   arrangements might otherwise make major contributions to it.  Looked
at
   differently, private arrangements tend to be examples of classic
   cream-skimming behavior:  the needs which can be most easily financed
are
   met, making it more difficult to meet needs -- possibly more serious
ones
   -- that are less readily financed.

 o Network arrangements set up for a particular project tend to collapse
when
   that project ends, leaving people who had learned to depend on a
certain
   level of communications and connectivity without it.

The decision to install a private network -- or, more often, simply a
private
star-type mail polling arrangement -- often results from lack of
understanding
```

```
or consideration of long-term implications.  Many of the people who are
normally considered experts can make poor guesses and give poor advice if
cost
and technology tradeoffs are radically different than they are in areas
with
established telecommunications and network infrastructures.  This
situation
prevails, almost by definition, in many less developed areas.  Although
the
reasons may be different, it is also prevalent in many areas of Eastern
Europe
and other portions of the former Soviet area of influence.  A very
similar
situation occurs in the US K-12 arena, where the cost of a single phone
line
can be a major administrative obstacle.

Even when networks develop within an area, without significant impetus
from
"outside", the user community that drives the installations may turn out
to be
the wrong one in the long term.  For example, the history of starting
networks
and network connections has predominantly rested in computer science and
computer technology-oriented departments, businesses, and other data
communities.  There is some history of these communities constructing
networks
for their own use and then using various mechanisms - costs, perceived
complexity, or lack of user support to then hoard the resource.

Then, when computer user communities - scientists, educators, or the
general
public - for whom the computer is a tool for communication or computation
(but
not an interesting device in its own right) need access to networking
technology, they often need to start over.  At the same time, those user
communities are much larger and, in many places, represent the largest
potential user community.  They may ultimately have access to greater
political, cultural, and financial resources than computer technologists.
```

```
At the same time, the community which is more computer-oriented can, and
typically has, managed to establish communications when they see it as
sufficiently important.  While they may not understand optimal approaches
for
a given area, they usually have access to sufficient information to make
something work.  Other scientists and organizations have often had less
success, since they do not have the technology readily available and may
believe (or "know") many things that are untrue, such as the need to
focus
computer networks around mainframe systems or centralized remotely-
accessed
hosts or the requirement for very high-bandwidth circuits to do anything
useful.  These misperceptions are encouraged by publicity releases from
vendors of high-end systems and by a media focus on the "next
generation".

These problems, and the behavior patterns that cause them, are not
limited to
international development or developing countries.  For example, in the
US K12
(kindergarten through 12th grade educational) community, few of the
innovative
teachers, the potential initial users who perceive the benefits in
advance,
have the political and economic power to affect their networking destiny.
Network inertia and data hoarding are rife in the administrative
infrastructure, often leading to either no connections at all or to
restricted
private single-function networks.  Theories of technological trickle-down
are
routinely cited but regularly disproven in the data networking arena.

All in all, it may have become a little bit too easy to set up a
"network".
Maximum effectiveness in network-building requires a different approach
in
which we phase out remote dialup arrangements in favor of "hosts" co-
located
with the users and where possible phase out -- and stop creating --
private
```

```
arrangements in favor of shared connections and infrastructure.

We need to improve our structure of information about who is interested
and
what is already operating in a particular area so that, if there is will
to
cooperate, every effort in a particular area reinforces every other
effort and
strengthens the links to the outside.  This, in turn, leads efforts down
the
path toward regional backbones as the most effective mechanism for
providing
adequate bandwidth through, and out of, the countries and regions.
Interested
parties can best leverage their own needs into effective networks if they
have
information about other interests and activities.  If existing sites are
not
inclined to cooperate with new ones, the best solution is to simply
develop
parallel infrastructure, gradually leaving them out.  That, of course,
requires the same databases, training, and information as would be the
case if
there were no existing connections to the area.

But, if we fail to move in ways that consolidate efforts and lead to
better
communications and interconnections within areas and between projects and
disciplines, we shall see increasing intercommunication and connection
difficulties among people and groups who "have email" or are "connected
to
networks".

------------------------------------
JOHN KLENSIN holds an S.B. and Ph.D. from MIT.  He is director of the
INFOODS
Secretariat for the United Nations University and was until recently
Principal
Research Scientist at MIT.  He has worked on or led major projects in
statistical and scientific database management, interchange of very
```

```
complex
data, information location and retrieval with uncertain classification,
data
analysis and modelling, and the impact and influence of communications.
He
has tried to use computer networks to enable applications and non-expert
users
since the early days of the ARPANet, and has occasionally succeeded.  He
is a
member of the Internet Engineering Task Force and chaired the recent
working
group on extensions to the SMTP protocol.  He is also a member of ACM and
immediate past chair of its Standards Committee, IEEE, the American
Statistical Association, and the International Association for
Statistical
Computing.  He can be reached as Klensin@INFOODS.UNU.EDU.

RANDY BUSH is a compiler netware, and tools hacker, and too often a
software
engineering manager.  Residing in Portland Oregon US, he is currently a
software architect at Olsen and Assoc., Zurich.  He has been a user and
occasional implementor of networking for a few decades, and is a member
of the
Modula-2 language committees and other lost causes.  He has been involved
in
in integration of appropriate networking technology in the developing
world
for over four years, using FidoNet, UUCP, and TCP/IP.  His email address
is
randy@psg.com.

Dr. Klensin and Mr. Bush are the principals in the Network Startup
Resource
Center, an effort to provide technical assistance to people developing
computer networks, especially low-end ones in developing areas.   That
effort,
and the development of this paper, are supported in part by NSF Grant No.
NCR-9216064.
```

CONTACT US

Network Startup Resource Center (NSRC)
1600 Millrace Drive, Suite E256
University of Oregon
Eugene, OR 97403-6232 USA
PHONE: +1-541-346-3547

FOLLOW US