

1993-07-21

**ISO/IEC JTC1/SC6
TELECOMMUNICATIONS AND INFORMATION
EXCHANGE BETWEEN SYSTEMS
Secretariat: U.S.A. (ANSI)**

Title: Proposed Technical Plan of Work for the ECFF Project

Source: USA (ANSI)

Project(s): JTC1.06.36

Status: Contribution to the SC 6 and SC 6 working group meetings in Seoul in October 1993.

Requested Action:

Attachment(s):

Distribution:

The purpose of this paper is to summarize the results of our discussion and analysis of the architectural basis for the ECFE work, and to recommend a program for applying these results to the ECFE problems. The architectural work stresses the following concepts:

- 1) *Three major phases of communication: enrollment, allocation, and data transfer.* These are always present in any communication, but in some cases they have been so simple that they have not been explicitly recognized: connection-oriented ignored enrollment; connectionless, enrollment and allocation. We have found that multicast requires careful attention to all three, and that explicitly recognizing all three in all forms of communication allows considerable unification and consistency in the architecture and in protocol design (see the US contribution entitled “The Architecture of Multicast Transport”).
- 2) *The separation of mechanism and policy.* This is not a new concept, but it has never been explicitly recognized in OSI protocol design. From our work, we have concluded that the concept of “context” that has been used in the upper layers can be applied to the definition of all protocols, and should be viewed as the means for negotiating the set of policies in force on an association.
- 3) *Elimination of the binary opposition of “connectionless” and “connection-oriented”* in favor of a more general continuum of which “connectionless” and “connection-oriented” are extreme end-points. A synthesis of connection-oriented and connectionless has been proposed in the US that allows them to be described as two extremes of a single property: the amount of state information shared among communicating entities.
- 4) *Recognition of a wide range of forms of multicast.* All of the proposals that have been presented to SC 6 to date have focussed on one form or another, but have failed to make it clear precisely which aspect of the problem they were attacking. This has made it difficult to compare the proposals, and has resulted in confusion over what each proposal is attempting to solve.

One of the most important things to come from this work was the realization that no single solution could cover the entire spectrum. Each choice of policies and mechanisms identifies a point in a multi-dimensional “solution space”, and each choice is valid for a particular set of criteria. Therefore, what we need is a way to tailor protocols to particular “operating regions” within the solution space.

With all of this in mind, let us consider each of the items in the list of topics described in the ECFE Guidelines (SC 6/N 7309).

Protocol Operation Streamlining

Multicast — From the work on extending TP4 for multicast described in the US contribution “Multicast Extensions to Class 4 of the Transport Protocol,” and from the architectural work described in the US contribution referred to in item 1 above, it is clear that the different forms of multicast Transport can be accommodated by defining and installing different sets of policies with respect to a common underlying set of Transport mechanisms. It is not necessary to define a completely separate new Transport protocol every time a new set of requirements for multicast (which can be expressed as a set of multicast policies) is recognized.

Reduced PDU Exchange — Most of these issues can also be accommodated by the definition of appropriate policies, to determine what allocation policy is used, how often acks are sent, how often credit or rate changes are made, etc. Experts in the US have worked out examples of such

policies for sequencing, allocation, and loss and duplicate detection mechanisms in enough detail to convince us that this approach would accommodate most of the possibilities being considered.

Selective Protocol Features — Virtually all of these features can be accommodated by having the ability to negotiate a specific set of policies at connection establishment time. At this point, we are not ready to suggest (for implementation reasons) that we negotiate mechanism as well. There is some concern that this might adversely affect the ability to create efficient implementations. However, it is certainly possible to negotiate the binary distinction between “no mechanisms,” i.e. a null policy, and mechanisms.

Large Sequence Number Space — There is not much in these topics that addresses this problem. We note only that this problem can be alleviated by the proper choice of the units in which sequence numbers are incremented, i.e. PDUs rather than octets, and by appropriate choice of PDU size.

Protocol Data Unit Encoding

Reduced PDU set — We have seen with TCP that reducing the PDU set to one is too few, and there are a lot of examples of having too many, but what is the rule? From the example mechanism specifications referred to above, it is clear that the rule should be that PDU types are closely associated with each mechanism: basically, one type of PDU per mechanism. While this should be kept in mind when designing a new protocol, this is definitely not a sufficient reason alone to redesign a protocol.

PDU Field Length and Placement — There is not much in the architectural work that addresses this issue. From the debate on the topic that has taken place within SC 6, it appears that there is general agreement that fields should be placed on 8 and 16 bit boundaries, that fields should be easy to find, that variable length fields should be avoided, and that if variable length fields are required, they should be placed at the end of the PCI. Beyond that it appears that other rules are dependent on the particular implementation strategy.

Efficient Checksum — This can be handled by allowing the policy for the data corruption mechanism to negotiate the checksum to be used. This would allow the checksum chosen to be appropriate to the characteristics of the data transfer and the error characteristics of the layer below.

Process Based Sequence — This issue is much more controversial than the others. There have been arguments that it makes a big difference, and counter arguments that it depends on the implementation strategy. As far as the checksum example that is often used, there seems to be as much hardware that would prefer the checksum at the beginning of the PDU as at the end. This argument seems to have more weight for protocols in the network layer and below where relaying is taking place. It seems very weak for a transport protocol where the PDU has been in memory for some time, and isn't being relayed but delivered to a local process. It would appear that as long as the checksum is easily accessible (e.g., aligned on a convenient byte boundary) most situations can be easily accommodated.

QoS — This has been a problem for years. There has been no shortage of parameters proposed for describing QoS, but it has been very difficult to determine how changes in these parameters are to affect the behavior of protocols. This has led to the following edict: If one can't relate a QoS parameter to how it affects a mechanism in some layer, it isn't a useful QoS parameter. However, one can use policy as the means to attack this problem. QoS parameters can be defined as abstractions of parameters that control the selection and expression of policies. We believe that this approach should form the basis for future work on QoS.

Recommendations

Currently, there are multiple proposals for the various pieces of the ECFF work. Some are proposing a protocol that is targeted at high speed and a portion of the multicast issues; others have been addressing particular aspects of multicast, arguing that the other issues can be dealt with separately. Considering all of this, we offer the following proposal for proceeding with the technical work on ECFF:

- 1) Develop an amendment to ISO/IEC 8072 incorporating a multicast transport service, as described by the New Project proposal from AFNOR contained in SC 6/N 8124 (on which the US has commented, in its NP ballot response, with respect to more carefully bounding the scope of the project).
- 2) Consider separating mechanism from policy in TP4 and adding an optional Context field. If the Context field is not present, then the default policies are the ones currently defined by TP4. If the context field is present, then the policies specified by the context are used.*
- 3) To support this extension to TP4, establish rules for defining Contexts and the policies that go with them and setting up the necessary registration authorities.

This provides a backward compatible path forward that will accommodate the various forms of multicast and other new ECFF capabilities that have been proposed by National Bodies to SC 6, with the exception of the PDU Field Length and Placement and the Process Based Sequence issues. These by their nature cannot be handled without sacrificing backward compatibility, which we believe would be a serious mistake. This would appear to be a good application of the “80-20” rule: 80% of the issues can be solved at a reasonable cost, but it would cost much more to solve the other 20%.

With this TP4*bis* in place, we could define policies that address the ECFF issues for general as well as specific operating environments. Groups could work on the issues of greatest importance to them knowing that solutions would all plug into a common base. Also, rather than having to evaluate whole new protocols that differ in several aspects, this approach would allow evaluation of changes to be restricted to essentially “one variable.” Perhaps the greatest advantage of this approach is that we avoid the overhead of defining and maintaining a large set of completely different protocols, accepting in its place the simpler problem of defining and managing sets of policies.

This proposal also has the advantage that it does not require complete agreement on strategies to address the various aspects of ECFF and multicast. With this approach, it is possible to accommodate not only the requirements of major applications, but also those of more specialized applications, which may not be seen as having major markets. For both situations one could develop policies (much less work than developing a protocol) to solve the various ECFF or multicast requirements while still using a common transport protocol with common mechanisms and PDU formats (and backward compatibility with the existing TP4 standard).

This lets everyone win. It also provides a powerful research vehicle that allows one to make controlled changes to a protocol and develop an understanding of the effects of that one change.

* The work on the upper layers of OSI has made extensive use of the concept of context in the presentation and application layers. Further work will be required to determine precisely how to apply this concept to work in the lower layers.