

Network Working Group  
Request for Comments: 3774  
Category: Informational

E. Davies, Ed.  
Nortel Networks  
May 2004

## IETF Problem Statement

### Status of this Memo

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

### Copyright Notice

Copyright (C) The Internet Society (2004). All Rights Reserved.

### Abstract

This memo summarizes perceived problems in the structure, function, and processes of the Internet Engineering Task Force (IETF). We are attempting to identify these problems, so that they can be addressed and corrected by the IETF community.

The problems have been digested and categorized from an extensive discussion which took place on the 'problem-statement' mailing list from November 2002 to September 2003. The problem list has been further analyzed in an attempt to determine the root causes at the heart of the perceived problems: The result will be used to guide the next stage of the process in the Problem Statement working group which is to recommend the structures and processes that will carry out the corrections.

## Table of Contents

1.	Introduction: Issues/Problems in the IETF Process . . . . .	2
1.1.	Consequences of Past Growth . . . . .	3
1.2.	The Aim is Improvement, not Finger-pointing . . . . .	4
1.3.	Perceived Problems - Consensus on Solutions . . . . .	4
2.	Root Cause Problems . . . . .	5
2.1.	Participants in the IETF do not have a Common Understanding of its Mission . . . . .	5
2.2.	The IETF does not Consistently use Effective Engineering Practices . . . . .	7
2.3.	The IETF has Difficulty Handling Large and/or Complex Problems . . . . .	9
2.4.	Three Stage Standards Hierarchy not properly Utilized . . . . .	11
2.5.	The IETF's Workload Exceeds the Number of Fully Engaged Participants . . . . .	12
2.5.1.	Lack of Formal Recognition . . . . .	13
2.6.	The IETF Management Structure is not Matched to the Current Size and Complexity of the IETF . . . . .	13
2.6.1.	Span of Authority . . . . .	13
2.6.2.	Workload of the IESG . . . . .	13
2.6.3.	Procedural Blockages . . . . .	15
2.6.4.	Consequences of Low Throughput in IESG . . . . .	15
2.6.5.	Avoidance of Procedural Ossification . . . . .	15
2.6.6.	Concentration of Influence in Too Few Hands . . . . .	16
2.6.7.	Excessive Reliance on Personal Relationships . . . . .	17
2.6.8.	Difficulty making Technical and Process Appeals . . . . .	18
2.7.	Working Group Dynamics can make Issue Closure Difficult . . . . .	18
2.8.	IETF Participants and Leaders are Inadequately Prepared for their Roles . . . . .	19
3.	Security Considerations . . . . .	20
4.	Acknowledgements . . . . .	20
5.	References . . . . .	21
5.1.	Normative References . . . . .	21
5.2.	Informative References . . . . .	21
6.	Editor's Address . . . . .	21
7.	Full Copyright Statement . . . . .	22

## 1. Introduction: Issues/Problems in the IETF Process

Discussion started in the second half of 2002 has shown that a significant number of problems are believed to exist in the way the Internet Engineering Taskforce (IETF) operates. Before attempting to change the IETF procedures and rules to deal with these problems, the IETF should have a clear, agreed-upon description of what problems we are trying to solve.

The Problem Statement working group was chartered to create this document, which contains a description of the problems, and to use this analysis to suggest processes to address the identified problems.

Taken in isolation, this document may appear to be exceedingly negative. The IETF needs to refresh its management and processes to address today's challenges, but it should not be forgotten that the IETF has produced a large body of high quality work which has lead to an extremely successful and pervasive network infrastructure. Against this background, we should see the current document as a necessary piece of self-criticism leading to renewal and continued success. The discussion of the positive aspects has been deliberately confined to the IETF Problem Resolution Processes document [5] which considers the core values that the IETF needs to maintain whilst correcting the problems that participants perceive as affecting the IETF at present.

The raw material for this document was derived by summarizing the extensive discussions which initially took place on the 'wgchairs' mailing list and subsequently on the 'problem-statement' mailing list from November 2002 through to September 2003, incorporating additional input from relevant drafts published during this period (see [2], [3] and [4]), and the minutes of recent plenary discussions. This produced a list of perceived problems which were classified into a number of related groups using a classification suggested by the processes which go on in the IETF.

This document has digested these perceived problems into a small set of root cause issues, and a short list of subsidiary issues which appear to be the most pressing items engendered by the root cause. This list is set out in Section 2.

Section 1.1 gives a short explanation of the thinking that has taken place in coming to the current view of the root causes.

The original summary of perceived problems has been posted to the Problem Statement Working Group mailing list so that it can be referred to in future. Note that it remains classified according the original scheme so that the raw data is available if alternative root cause analysis is needed.

### 1.1. Consequences of Past Growth

As the problems of the IETF were examined, it became clear that they are neither new nor are they symptoms of a problem which is novel in the science of organizations.

The IETF started off as a small, focused organization with a clearly defined mission and participants who had been working in this area for a significant period of time. Over the period 1989-1999, the IETF grew by a factor of ten or more in terms of number of participants, and volume of work in progress. The effects of this growth have been compounded by the extension of the scope of the IETF which makes the work much more varied. Also during this period, the Internet has become more complex and the requirements placed on it by a far larger user community have changed as the network has come to have a pivotal role in many areas of life.

Many of the problems and symptoms appear to be fundamentally caused by the organization failing to fully adapt its management structure and processes to its new larger size and the increased complexity of the work. The IETF has also failed to clearly define its future mission now that the initial mission has been completed or outgrown.

These failures are just those that afflict many small organizations trying to make the transition from a small organization, which can be run informally and where essentially all participants fully share the aims, values, and motivations of the leadership, to a medium sized organization, where there are too many participants for informal leadership and later arrivals either do not fully understand or have a different perception of the ethos of the organization.

Some IETF participants have been aware of these issues for a long time. Records dating back to at least 1992 drew similar conclusions.

#### 1.2. The Aim is Improvement, not Finger-pointing

Many of the problems identified in this memo have been remarkably persistent over a 15-year period, surviving a number of changes in personnel. We see them as structural problems, not personnel problems. Blame for any of the perceived problems should not be directed to any individual. The sole aim of this review process is to identify how the IETF can improve itself so that it knows what it is about and becomes fit for that purpose in the shortest possible time frame.

#### 1.3. Perceived Problems - Consensus on Solutions

The working group participants emphasize that both the long list of problems and the root cause issues that were derived from them are problems that are believed to exist by a significant constituency, either on the mailing list and/or in private discussions. We also note that many of these problems appear to be of long standing, as a

very similar list has survived from the discussions in the first POISED working group that took place prior to the IETF organizational changes approved in 1992.

We, in line with many contributors to the mailing list, believe that it is important to try and identify what appear to be the root causes of the perceived problems, but trying to prioritize or assign a relative importance to the problems would not be useful: rough consensus on an unordered list of real and important root causes will be sufficient. The root causes identified will provide a guide in setting up the processes needed to resolve the problems: the perceived problems can be viewed as multiple symptoms of the root causes which should provide input to those trying to resolve the problems in achieving consensus on solutions.

## 2. Root Cause Problems

This section forms the heart of this analysis, and lists the issues which we believe lie at the core of the problems. Apart from the first issue which is fundamental, the problems are not necessarily in priority order, but they will be seen to be interlinked in various ways.

### 2.1. Participants in the IETF do not have a Common Understanding of its Mission

The IETF lacks a clearly defined and commonly understood Mission: as a result, the goals and priorities for the IETF as a whole and any Working Groups (WGs) that are chartered are also unclear.

The IETF needs to understand its mission in the context of the greatly increased scope and complexity of the Internet, and the changing requirements of the much larger user community that the success of its previous work has engendered.

The lack of a common mission has many consequences, of which the principal ones appear to be:

- o The IETF is unsure what it is trying to achieve and hence cannot know what its optimum internal organization should be to achieve its aims.
- o The IETF cannot determine what its 'scope' should be, and hence cannot decide whether a piece of proposed work is either in-scope or out-of-scope.
- o The IETF is unsure who its stakeholders are. Consequently, certain groups of stakeholder, who could otherwise provide

important input to the process, have been more or less sidelined because it has seemed to these stakeholders that the organization does not give due weight to their input.

- o Working Groups can potentially be hijacked by sectional interests to the detriment of the IETF's mission.
- o The misty vision has inhibited the development of roadmaps that would inform the IETF's stakeholders of our longer term intentions, as well as restricting the associated architectural views to an outline top level view which does not fully reflect the developing nature of the Internet. It would be desirable to have roadmaps and architectural views for portions of work which extend beyond a single working group: it may also be the case that it is no longer possible to fit the whole Internet within a single architecture.
- o The IETF is unable to determine explicitly what effect it desires to have in the marketplace, and is therefore unable to determine what requirements of timeliness are appropriate when planning work and setting expectations for stakeholders which will further the IETF's mission.
- o The lack of precision regarding our goals leads to WG charters and requirements that are poorly thought out and/or not aligned with the overall architecture. The resulting poorly defined charters are a major factor in poor quality and/or late deliveries from some WGs and the total failure of other WGs.
- o The IETF needs to avoid focusing on a too-narrow scope of technology because this would be likely to blinker the IETF's view of 'the good of the Internet', and will harm the long-term goal of making the Internet useful to the greatest number stakeholders; this argues for allowing a relatively wide range of topics to be worked on in the IETF - cross-fertilization has always been one of the IETF's strengths.

An additional barrier to achieving a common understanding is that the IETF does not have a recognized forum in which all stakeholders participate and in which organization wide consensus might be reached. Plenary meetings during regular IETF meetings allow a large cross-section of the community to offer views, but there is not generally sufficient time to achieve consensus and there is no single mailing list which all stakeholders can be guaranteed to monitor.

The IETF creates standards and is therefore necessarily a Standards Development Organization (SDO), but many participants would like to differentiate the IETF and its way of working from the 'conventional'

SDOs which emphasize corporate involvement and mandated delegates. Externally, the IETF is often classified with these conventional SDOs, especially by detractors, because the differentiation in the IETF's mission and processes and the rationale for those differences are not clear. This can lead to the IETF being misunderstood by other SDOs which can make communications between SDOs less effective, harming the IETF's ability to achieve its mission.

## 2.2. The IETF does not Consistently use Effective Engineering Practices

For an organization with 'engineering' in its title and participants who are likely to trot out the statement "Trust me, I'm an engineer!" when confronted with the need to find a solution to a particularly knotty problem, the IETF has, at least in some cases, extremely ineffective engineering practices. Effective engineering practices, as used here, covers both the techniques used to derive and verify the technical solutions needed, and the management and organizational strategies that are commonly accepted to help with the engineering process.

A major symptom of this lack is that WGs do not consistently produce timely, high-quality, and predictable output. As discussed in Section 2.1, this problem is exacerbated because the IETF currently finds it difficult to determine what is timely, and hence what are appropriate deadlines for the delivery of WG output. Some of the contributing problems which interfere with effective engineering in WGs include:

- o Failure to ensure that there is a uniform view in the WG of the scope of the WG activity, especially the intended purpose of the solution.
- o Failure to identify the issues that need to be resolved at an early stage (before the design is frozen), and/or then to ensure that there is a uniform view in the WG of the issues that need to be resolved to bring the work to a satisfactory conclusion.
- o Failure to identify and articulate engineering trade-offs that may be needed to meet the deadlines that the WG has set without inappropriately reducing the 'fitness for purpose' for the intended customers.
- o Continued refinement of the solution beyond the point at which it is adequate to meet the requirements placed on it by the intended purpose.

The IETF standards engineering process is not set up to deliver iterative process improvement. Particular areas that need improvement include:

- o The charter may not be sufficiently detailed to document the process and timeline to be followed by the WG. Additional documents may be needed, such as a roadmap or detailed plans.
- o Poorly defined success criteria for WGs and individual documents.
- o Lack of written guidelines or templates for the content of documents (as opposed to the overall layout) and matching lists of review criteria designed to achieve appropriate quality in output.
- o Lack of auditing against explicit criteria throughout the standards development process.
- o Lack of review, especially early review, by reviewers who are not directly interested members of the WG, and by subject matter experts for topics related to, but not necessarily the immediate focus of the document.
- o Lack of documentation about likely problem areas that might arise due to interactions with other popular IETF protocols.
- o Lack of metrics to measure the achievement of the desired quality and the performance of both WGs and the whole IETF.
- o Lack of metrics and 'post mortem' procedures to drive the improvement of the standards development and other IETF processes.
- o Lack of criteria for determining when a piece of work is overrunning and/or is unlikely to be concluded successfully, either at all or within an acceptable time frame. Lack of process for extending the time frame, adjusting the scope, or terminating the work item or the whole Working Group.
- o Automated tools to support the engineering process are minimal.
- o Despite its commitment to 'running code', the IETF is not proactive in providing ways for developers to verify their implementations of IETF standards.

In addition, IETF processes, and Working Group processes in particular, suffer because commonly accepted Project Management techniques are not regularly applied to the progress of work in the organization.



- o Project entry, goal setting, dependency identification, coordination, and tracking processes are all either missing or implemented less effectively than the norm for commercial organizations in related activities. Dependencies and coordination should cover both other WGs within the IETF and any outside SDO with which the IETF is collaborating.
- o Charters regularly fail to set enough milestones with sufficiently small granularity at which progress of WGs, individuals, and documents can be evaluated. Also, WGs often do not make more detailed work plans to refine the charter plans.
- o The acceptable deadlines for finishing a piece of work, and the criteria used to determine them, are rarely, if ever, documented. Also, the estimated time required to complete the work often differs widely from the time actually taken. The combination of these factors makes determining the feasibility of delivering within the required time frame, and then adjusting the scope of the work to fit the time frame requirements, extremely difficult.

One problem which the IETF does not appear to suffer from is excessive bureaucracy, in the sense that transfer of information is generally kept to the minimum necessary to accomplish the task. It is important that any changes introduced do not significantly increase the bureaucratic load whilst still recording sufficient information to allow process improvement.

Finally, even where the IETF does have Engineering Practices defined, there are frequently cases where they are ignored or distorted. One area of particular concern is the tendency for protocols to be assessed and issues resolved primarily through static analysis of the written specification rather than by practical experiment with 'running code'.

### 2.3. The IETF has Difficulty Handling Large and/or Complex Problems

The IETF has historically been most successful when dealing with tightly focused problems that have few interactions with other parts of the total problem solution. Given that the Internet has become more complex, such tightly focused problems are becoming the exception. The IETF does not always seem to be aware of the interactions between protocols that are bound to be thrown up by deployment in more complex situations and so fails to minimize the chances of unwelcome consequences arising unforeseen when a new protocol is deployed. This may be exacerbated by inadequate review from outside the WG as suggested in Section 2.2.

IETF standardization procedures are optimized for tightly constrained working groups and are generally less effective if 'engineering in the large' is needed to reach a satisfactory solution. Engineering in the large can encompass many aspects of system design including:

- Architecture
- Frameworks
- Security
- Internationalization

The IETF has historically standardized protocol components rather than complete systems, but as we have learned more about the ways in which systems on the Internet interact, design of components needs to take into account more and more external constraints, and the understanding of these constraints tends to require more engineering in the large.

Part of the cause of this difficulty may be that the formal reporting structure of the IETF emphasizes communication between the Internet Engineering Steering Group (IESG) through the ADs and the WGs, and does not place much reliance on inter-WG communications:

- o The IETF is not consistently effective at resolving issues that cross WG or area boundaries.
- o The IETF does not possess effective formal mechanisms for inter-WG cooperation, coordination, or communication, including the handling of dependencies between deliverables and processes specified in WG charters.
- o The IETF does not have an effective means for defining architectures and frameworks that will shape the work of multiple WGs.

The IETF also has to work with other SDOs, and the liaison mechanisms for coordination and cooperation do not always work efficiently. This needs to be remedied because some of the interactions which IETF work has to take into account will involve protocols and systems standardized by these other SDOs.

A possible consequence of the need for more engineering in the large is that protocol specifications have become larger: as a result they now take longer to develop. Some people perceive that this is because the IESG has tended to require protocol specifications to specify an entire system, instead of simple component protocols, leading to feature bloat and applicability only to a narrow range of applications (see also Section 2.4). On the other hand, others believe that the IESG has approved simple component protocols without

an adequate understanding of the systems and contexts in which the protocols might be used. These problems appear to be two additional aspects of the general problem that the IETF has with handling large and/or complex systems.

#### 2.4. Three Stage Standards Hierarchy not properly Utilized

The current hierarchy of Proposed, Draft, and Full Standard maturity levels for specifications is no longer being used in the way that was envisioned when the stratification was originally proposed. In practice, the IETF currently has a one-step standards process that subverts the IETF's preference for demonstrating effectiveness through running code in multiple interoperable implementations. This compresses the process that previously allowed specifications to mature as experience was gained with actual implementations:

- o Relatively few specifications are now progressed beyond Proposed Standard (PS) to Draft Standard (DS) level, and even fewer to Full Standard (FS).
- o It is widely perceived that the IESG has 'raised the (quality) bar' that standards have to pass to be accorded a PS status. Protocol developers may be required to specify a complete system rather than an interface in order for their specification to be approved as a PS (see also Section 2.3).
- o In spite of the apparently higher quality hurdle, implementation or deployment experience is still not required, so the IETF's guiding principle of 'rough consensus and running code' has less of a chance to be effective.
- o There appears to be a vicious circle in operation where vendors tend to deploy protocols that have reached PS as if they were ready for full production, rather than accepting that standards at the PS level are still under development and could be expected to be altered after feature, performance, and interoperability tests in limited pilot installations, as was originally intended. The enthusiasm of vendors to achieve a rapid time to market seems to have encouraged the IETF in general and the IESG in particular to attempt to ensure that specifications at PS are ready for prime time, and that subsequent modifications will be minimal as it progresses to DS and FS, assuming effort can be found to create the necessary applicability and interoperability reports that are needed.
- o The three stage hierarchy is, accordingly, seen to be excessive.

- o There is no formal bug reporting or tracking system in place for IETF specifications.
- o The periodic review of protocols at PS and DS levels specified in [1] are not being carried out, allowing protocols to persist in these lower maturity levels for extended periods of time, whereas the process would normally expect them to progress or be relegated to Historic status.
- o No individual or body is given the task of 'maintaining' a specification after the original WG has closed down. Specifications are generally only updated when a need for a new version is perceived. No attempt is normally made to correct bugs in the specification (whether they affect operation or not) and the specification is not updated to reflect parts of the specification that have fallen into disuse or were, in fact, never implemented. This is, in part, because the current procedures would require a standard to revert to the PS maturity level, even when specification maintenance is carried out. This occurs even if the changes can be demonstrated to have no or minimal effect on an existing protocol at the DS or FS level.

#### 2.5. The IETF's Workload Exceeds the Number of Fully Engaged Participants

There are a number of respects in which IETF participants and contributors appear to have become less fully engaged with the IETF processes, for example:

- o Although there may be large attendance at many WG meetings, in many cases, 5% or less of the participants have read the drafts under discussion or that have a bearing on the decisions to be made.
- o Commitments to write, edit, or review a document are not carried out in a timely fashion.
- o Little or no response is seen when a request for 'last-call' review is issued, either at WG or IETF level.

This might be because contributors have less time available in their work schedule during the downturn of the Internet business climate between 2001 and 2003. Yet, this is not the whole story, as there were signs of this effect back at the height of the Internet's boom in 2000.

This problem exacerbates the problems the IETF has had with timely delivery and may weaken the authority of IETF specifications if decisions are seen to be taken by badly informed participants and without widespread review.

#### 2.5.1. Lack of Formal Recognition

Beyond RFC Authorship, WG Chair positions, Directorate positions, or IESG and Internet Architecture Board (IAB) membership, the IETF does not offer formal recognition of contributions to the IETF. This potentially acts as a disincentive to continued engagement and can lead to useful and effective participants leaving because they cannot obtain any recognition (the only currency the IETF has to pay participants), which they use to fuel their own enthusiasm and help justify their continued attendance at IETF meetings to cost constrained employers. Note: Using Leadership positions as rewards for good work would probably be damaging to the IETF. This paragraph is meant to indicate the need for other types of rewards.

#### 2.6. The IETF Management Structure is not Matched to the Current Size and Complexity of the IETF

The management and technical review processes currently in place were adequate for the older, smaller IETF, but are apparently not scalable to the current size of the organization. The form of the organization has not been significantly modified since 1992, since when the organization has undergone considerable further growth. The scope of IETF activities has also been extended as the Internet has become more complex.

##### 2.6.1. Span of Authority

Overt authority in the IETF is concentrated in the small number of people sitting on the IESG at that time. Existing IETF processes work to funnel tasks on to this small number of people (primarily the Area Directors (ADs) in the IESG). This concentration slows process and puts a very large load of responsibility on the shoulders of these people who are required to act as the senior management for Working Group (WG) chairs, as well as acting as quality backstops for the large number of documents issued by the IETF. The situation has not been helped by the widening of the scope of the IETF, which has resulted in somewhat more WGs and a need for a very broad spectrum of knowledge within the set of ADs.

### 2.6.2. Workload of the IESG

With the current structure of the IETF and IESG, the workload imposed on each of the ADs is almost certainly well beyond the capabilities of a single person.

The current job description for an AD encompasses at least the following tasks:

- o Interacting with WGs
- o Understanding network and computer technology in general, and their own area in detail
- o Cross-pollinating between groups
- o Coordinating with other areas
- o Potentially, managing their Area Directorate team
- o Effectively providing technical management, people-management, and project supervision for their WGs
- o Reading (or at least skimming) every formal document which the IETF produces, and having an opinion on all of them, as well as all the Internet Drafts produced by the WGs in the area, and understanding the interactions between all these specifications.

Given the number of WGs which are now active, the increasing complexity of both the work being undertaken and the technology in general, together with the volume of documents being produced, makes it clear that only superhumans can be expected to do this job well. To make matters worse, these tasks are, in theory, a 'part time' occupation. ADs will normally have a conventional job, with the IETF activities as just one part of their job specification. This view has been reinforced by recent resignations from the IESG, citing the size of the workload as a primary factor. The IETF also has no mechanisms to nominate a temporary replacement or an assistant should an AD be incapacitated wholly or partially for a period.

The malign effects of this overload include:

- o Wear on the IESG: The IESG members are overworked which is bad for their health, humor, and home life, and may also result in conflicts with their employers if the IETF work impacts the IESG member's performance of their 'day job'.

- o Unhappiness in the IETF: IETF stakeholders perceive that IESG members are responding slowly, are not fully up-to-date with their technology, fail to pro-actively manage problems in their WGs, and are unable to keep communication channels with other groups open.
- o Recruiting shrinkage: The number of people who can imagine taking on an IESG post is steadily decreasing. It is largely limited to people who work for large companies who can afford to send IESG members to the IETF for the duration of their appointments. In the current business climate, fewer companies are able to justify the preemption of an important engineering and business resource for a significant period of time, and are more likely to put forward 'standards professionals' than their best engineers.

#### 2.6.3. Procedural Blockages

The current procedural rules combined with the management and quality roles of the ADs can lead to situations where WGs or document authors believe that one or two ADs are deliberately blocking the progress of a WG document without good reason or public justification. Appeal processes in these circumstances are limited and the only sanction that could be applied to the relevant ADs is recall, which has almost always been seen to be out of scale with the apparent offense and hence almost never invoked. This perception of invulnerability has led to a view that the IESG in general and the ADs in particular are insufficiently accountable for their actions to their WGs and the IETF at large, although the recent introduction of the Internet Draft Tracker tool makes it easier to determine if and how a document has become blocked, and hence to take appropriate steps to release it.

#### 2.6.4. Consequences of Low Throughput in IESG

If documents are inappropriately (or even accidentally) delayed or blocked as a result of IESG (in)action, this can cause much frustration inside the organization, a perception of disunity seen from outside the organization, and delay of standards, possibly to the point where they are too late to match market requirements: work which has been properly authorized as being within the scope of the IETF and properly quality checked during development, should almost never come up against such a blockage.

Delay in authorizing a BOF or chartering a new WG can delay the start of the process with similar effects.

It also appears that IESG delays are sometimes used to excuse what is actually slow work in WGs.

#### 2.6.5. Avoidance of Procedural Ossification

The systems and processes used by the IETF are generally designed around having firm general principles and considerable IESG discretion within those principles. It appears that the IETF is showing a disturbing tendency to turn IESG 'rules of convenience' into rigid strictures that cannot be violated or deviated from.

Up to now, IETF discussions of procedures have been driven by a model in which the procedural BCPs construct a framework for doing work, but the details of the framework are left for the IESG to fill in. When issues or crises have arisen, the IETF has generally avoided making specific procedural changes to compensate, instead realizing that we could not anticipate all cases and that 'fighting the last war' is not a good way to proceed.

This can only continue to work if the participants continue to trust the IESG to act fairly in filling in the details and making appropriate exceptions, without a great deal of debate, when it is clearly desirable. At present, the IETF appears to have lost sight of this flexibility, and is entangling itself in procedures that evolve from organizational conveniences into encumbrances.

#### 2.6.6. Concentration of Influence in Too Few Hands

Until the last couple of years, successive IETF Nominating Committees have chosen to give heavy weighting to continuity of IESG and IAB membership. Thus, the IETF appeared to have created an affinity group system which tended to re-select the same leaders from a limited pool of people who had proved competent and committed in the past.

Members of this affinity group tend to talk more freely to each other and former members of the affinity group - this may be because the affinity group has also come to share a cultural outlook which matches the dominant cultural ethos of the IETF (North American, English speaking). Newcomers to the organization and others outside the affinity group are reluctant to challenge the apparent authority of the extended affinity group during debates and consequently influence remains concentrated in a relatively small group of people.

This reluctance may also be exacerbated if participants come from a different cultural background than the dominant one. Such participants also tend to find it more difficult to follow the rapid and colloquial speaking style of native English speakers, and may consequently be effectively excluded from the discussion, even if maximum assistance is available by such means as real time Jabber logs and extensive text on presentation slides. Even on mailing



lists, people from other cultures may be reluctant to be as forthright as is often the case in discussions between North Americans; also, a person whose first language is not English may be daunted by the volume of mail that can occur on some mailing lists and the use of colloquialisms or euphemisms may cause misunderstandings if correspondents are not aware of the problem.

A further instance of the problems of concentration of influence potentially occurs when, from time to time, ADs have acted as WG chairs: conflict of interest might well arise in discussions between the IESG and any WG with an AD as its chair. Whilst care is usually taken to have a newly selected AD vacate any WG chair positions which might be held in his or her own area, the conflict can arise on the occasions when an AD has been used as the chair of a WG because it is clearly the right (or only possible) solution for the WG from an engineering and know-how position. Furthermore, given the known problem of workload for IESG members, there must be doubts as to whether an AD can or ought to be taking on this extra load.

#### 2.6.7. Excessive Reliance on Personal Relationships

The IETF is an intensely personal and individualistic organization. Its fundamental structure is based on individuals as actors, rather than countries, organizations, or companies as in most other SDOs.

This is also reflected in how the IETF gets its work done: the NOMCOM process, the WG Chair selection processes, and the activities of WGs are all reliant on personal knowledge of the capabilities of other individuals and an understanding built on experience of what they can be expected to deliver, given that there are almost no sanctions that can be applied beyond not asking them to do a similar task again. The relationship works best when it is two way - the person being asked to perform a task needs to be able to rely on the behavior of the person doing the asking.

In essence, the IETF is built on a particular kind of one-to-one personal trust relationship. This is a very powerful model but it does not scale well because this trust is not transitive. Just because you trust one person, it does not mean that you trust (i.e., know the capabilities of and can rely on) all the people that person trusts in turn.

The disruption caused when one set of relationships has to be replaced by another is clearest when an AD is replaced. The IETF does not keep personnel records or written plans, and formal process documentation is very sparse, so that incoming ADs have little information on which to base new relationships with WG chairs or Directorate members not already known to them.

A new AD has to build or bring along his or her set of trusted individuals. The AD will tend to prefer individuals from this set as WG chairs, unless there is a suitable outsider who was part of the team that brought the WG idea to the IETF. This tends to limit the AD's field of choice, particularly when asking for a 'stabilizing', 'advising', or 'process' chair to work with an enthusiastic newcomer in a difficult area. A breakdown of an established relationship (such as between an AD and a WG chair) can be very damaging to the work of the IETF, and it may not be immediately obvious to outsiders.

Another consequence of the reliance on personal relationships is that the IETF has very little institutional 'memory' outside the memories of the people in the process at a given time. This makes it more likely that failures will be repeated and makes process improvement more difficult (see Section 2.2).

#### 2.6.8. Difficulty making Technical and Process Appeals

When an individual thinks that the process has produced a result that is harmful to the Internet or thinks that IETF processes have not been adhered to, there is no mechanism to aid that individual in seeking to change that result.

#### 2.7. Working Group Dynamics can make Issue Closure Difficult

The IETF appears to be poor at making timely and reasonable decisions that can be guaranteed to be adhered to during the remainder of a process or until shown to be incorrect.

The problems documented in this section are probably consequences of the non-hierarchical organization of the IETF and the volunteer status of most participants. The enforcement measures available in a more conventional hierarchical corporate environment are mostly not available here, and it is unlikely that application of some well-known procedure or practice will fix these problems.

Participants are frequently allowed to re-open previously closed issues just to replay parts of the previous discussion without introducing new material. This may be either because the decision has not been clearly documented, or it may be a maneuver to try to get a decision changed because the participant did not concur with the consensus originally. In either case, revisiting decisions stops the process from moving forward, and in the worst cases, can completely derail a working group. On the other hand, the decision making process must allow discussions to be re-opened if significant new information comes to light or additional experience is gained which appears to justify alternative conclusions for a closed issue.

One cause that can lead to legitimate attempts to re-open an apparently closed issue is the occurrence of 'consensus by exhaustion'. The consensus process can be subverted by off-topic or overly dogmatic mail storms which can lead to the exclusion of knowledgeable participants who are unable to devote the time needed to counter the mail storm. The consequence may be an unrepresentative and unsatisfactory consensus which will tend to be re-opened, often leading to repeat discussions. Mailing lists, which are at the heart of the IETF WG process, are becoming increasingly ineffective at resolving issues and achieving consensus because of this phenomenon.

A single vocal individual or small group can be a particular challenge to WG progress and the authority of the chair. The IETF does not have a strategy for dealing effectively with an individual who is inhibiting progress, whilst ensuring that an individual who has a genuine reason for revisiting a decision is allowed to get his or her point across.

## 2.8. IETF Participants and Leaders are Inadequately Prepared for their Roles

Participants and leaders at all levels in the IETF need to be taught the principles of the organization (Mission and Architecture(s)) and trained in carrying out the processes, which they have to use in developing specifications, etc.

Part of the reason for the lack of training in the principles of the organization is that there is not currently an explicit formulation of these principles that is generally agreed upon by all stakeholders. Section 2.1 identifies that this shortage is a major problem.

The IETF currently has voluntary and inconsistent processes for educating its participants, which may be why significant numbers of participants seem to fail to conform to the proper principles when working in the IETF context.

The people in authority have generally been steeped in the principles of the IETF (as they see them) and first-time non-compliance by newer participants is sometimes treated as an opportunity for abuse rather than recognition of a training failure.

The IETF culture of openness also tends to tolerate participants who, whilst understanding the principles of the IETF, disagree with them and actively ignore them. This can be confusing for newer participants, but they need to be made aware that the IETF does not exclude such people. The IETF does not currently have a strategy for

dealing with the conflicts that can result from participants who disagree with the principles of the organization.

Lack of training, compounded with the perceived concentration of influence in the affinity group documented in Section 2.6.6, can lead to newcomers being ignored during discussions, consequently being ineffective, either in their own eyes or their employers. This may result in their departure from the IETF.

In addition, some participants are not aware of the problems that participants, who do not have English as their first language, may have with rapid speaking and the use of colloquialisms in both spoken and written communication. They are also not always aware of the possible cultural nuances that may make full participation more difficult for those who do not share the same outlook.

### 3. Security Considerations

This document does not, of itself, have security implications, but it may have identified problems which raise security considerations for other work. Any such implications should be considered in the companion document which will be produced setting out how the IETF should set about solving the identified problems.

### 4. Acknowledgements

Apart from the contributions of all those who provided input on the problem statement mailing list, the final reduction of the problems was especially assisted by the following people:

Rob Austein <sra@hacrn.net>  
Marc Blanchet <Marc.Blanchet@hexago.com>  
Dave Crocker <dcrocker@brandenburg.com>  
Spencer Dawkins <spencer@mcsr-labs.org>  
Avri Doria <avri@psg.com> (WG co-chair)  
Jeanette Hoffmann <jeanette@wz-berlin.de>  
Melinda Shore <mshore@cisco.com> (WG co-chair)  
Margaret Wasserman <margaret@thingmagic.com>

Special thanks are due to Margaret Wasserman for extensive reviewing of and contributions to the wording of Section 2.

The early part of the reduction of the problem statement mailing list input was done by Harald Alvestrand and the latter part by Elwyn Davies and the team acknowledged above. In total, there were approximately 750 extensive and thoughtful contributions (some making

several points). The thread was started by a call for volunteers in helping draft a problem statement, but quickly turned into a discussion of what the problems were.

In addition to the editorial team, the following people have provided additional input and useful feedback on earlier versions of this document: Harald Alvestrand, Randy Bush, Brian Carpenter, James Kempf, John Klensin, John Loughney, Keith Moore.

## 5. References

### 5.1. Normative References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

### 5.2. Informative References

- [2] Huston, G. and M. Rose, "A Proposal to Improve IETF Productivity", Work in Progress.
- [3] Blanchet, M., "Suggestions to Streamline the IETF Process", Work in Progress.
- [4] Hardie, T., "Working Groups and their Stuckees", Work in Progress.
- [5] Davies, E. and J. Hofmann, Eds., "IETF Problem Resolution Processes", Work in Progress.

## 6. Editor's Address

Elwyn B. Davies  
Nortel Networks  
Harlow Laboratories  
London Road  
Harlow, Essex CM17 9NA  
UK

Phone: +44 1279 405 498  
EMail: elwynd@nortelnetworks.com

## 7. Full Copyright Statement

Copyright (C) The Internet Society (2004). This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

## Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at [ietf-ipr@ietf.org](mailto:ietf-ipr@ietf.org).

## Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

