



« A Joint Effort of the [INTERNET MULTICASTING SERVICE](#) and [INTERNET SOFTWARE CONSORTIUM](#) »

[TRANSMITTAL](#) »[FITNESS DISCLOSURE](#) »[VIEW OUR PRESENTATION](#) »[PROPOSAL](#) »[SUPPLEMENTARY MATERIALS](#) »[SPREAD THE DOT](#) »[CONFIDENTIAL INFORMATION](#) »[SUPPLEMENTAL QUESTIONS](#) »[FINAL RESPONSE](#) »

Response to Final Evaluation Report

Table of Contents

- [1.](#) Preamble
- [2.](#) Performance
- [3.](#) Price
- [4.](#) Process
- [5.](#) Parable
- [6.](#) Purpose
- [§](#) References
- [§](#) Authors' Addresses
- [A.](#) Word Count
- [B.](#) Excerpts from the Adams NIC Proposal, 5/1992
- [C.](#) Document Formats

[DOC](#)

1. Preamble

[TOC](#)

The [Internet Multicasting Service](#) and the [Internet Software Consortium](#) are pleased to address these comments to the ICANN board in support of our bid [\[1\]](#) to operate the .org Top Level Domain (TLD) as provided for by the Seventh Update [\[2\]](#) and Final Evaluation Report. [\[3\]](#)

On June 26, 2002 we had the pleasure of addressing you in Bucharest. [\[4\]](#) We stressed five aspects of our proposal:

- Performance
- Price
- Process
- Parable
- Purpose

We would like to revisit each of these five aspects of our proposal, keeping in mind the following points:

- The IMS/ISC team has a widely recognized demonstrated ability to operate registries of significant scale.
- The IMS/ISC bid is very different from the bids that ICANN staff deemed "qualified," all of which have a similar

operational style and pricing philosophy as VeriSign.

- The reassignment of the .org TLD is a significant event in the transformation of the DNS. Higher prices from barriers to entry could well cost consumers US\$268 million over the next five years.

2. Performance

[TOC](#)

The IMS/ISC proposal was disqualified under Criterion 1 for not having shown a "demonstrated ability" to operate a TLD registry of "significant scale." (To be fair, the IMS/ISC proposal was not explicitly disqualified by staff, it was rated of lower quality than those bids that were considered "marginal," which in turn were rated lower than bids that were "highly qualified.") This disqualification was explained by staff in their Final Evaluation:

"There are those who will believe that only large registries could be competitive in this bidding process. Given the primacy of transferring a complex and large registry, .org, in a demonstrably stable manner, there must be some truth to this. This was well articulated in the RFP itself. In fact, staff was surprised by the number of proposals received, a number that exceeded initial expectations (hence the subsequent refund) given the focus of the RFP on 'demonstrated ability to operate a TLD of significant scale.'

This does not mean that all future ICANN gTLD awards will contain the same requirement. After all, none of those featured in this Report for further consideration started "large": ISOC's partner Afilias, NeuStar, and GNR all started from scratch last year. However, in this case, the Board is considering transfer of a large functioning registry, .org, containing well over 2 million registrants, and cannot responsibly risk jeopardizing stability through venturing too far into the unknown and the unknowable. This is the compelling fact for this reassignment. But there [sic] will likely be other opportunities for those who did not succeed in this particular award." [3]

The IMS/ISC team has a widely-recognized "demonstrated ability," including long experience producing [BIND](#) and [DHCP](#) software and operating nameservers for the root and TLDs, a [fully operational registration system](#) loaded with real .org data, and a [long track record of relevant experience](#) in all aspects of registry operation. Specifically:

1. We already run closely related services that are a key part of the Internet infrastructure, such as the [F Root Server](#) and are a significant player in TLD hosting. These services run on our established infrastructure and operate with the highest standards of availability, security, and responsiveness. A large part of provisioning a registry service that can scale is in the hosting and DNS distribution infrastructure. A careful examination of the hosting, transit, security, and other "bit-pushing" pieces of each bidder's infrastructure is a crucial part of any thorough evaluation.
2. Our OpenReg system provides a live, fully operational mirror of the .org zone, with fully operational support for an OT&E testbed, interoperability testing for protocol compliance, and extensive load and availability testing. Several ICANN-certified registrars have tested this system and found it identical in operation to, albeit considerably faster than, the current VeriSign-operated .org registry. Any Internet user can: [test a .org domain name](#) in our live mirror; [add sample data](#) to the OT&E testbed to test the RRP, Whois, and DNS components of our system; and, [examine performance and availability](#) of the system.
3. Our team members have [direct and highly relevant experience](#) running all aspects of registry services, including Whois services, DNS services, registrar services, and registry services. Our experience in this field dates back to the original registry services provided by SRI and our team members have worked with, worked for, and supported every major player, including ICANN.

In the Final Evaluation, staff indicated that new entrants to the market should cut their teeth on something less important than .org. They stressed that real experience running a real registry was the only way stability goals could be met. They then proceeded to list the amount of experience each of the "qualified" applicants had:

Applicant	SRS Experience	Whois Experience
=====	=====	=====
Dot Org Foundation	up to 16.5 Months	8.5 Months
GNR	0 Months	5 Months
Internet Society	10 Months	9 Months
NeuStar	7.25 Months	8.5 Months
RegisterOrg	up to 16.5 Months	8.5 Months
Unity	0 Months	5-6 Months
UIA	38 Months	11 Years

[Note: DNS Experience Not Tabulated.]

The registry "industry" is made up of three groups:

1. VeriSign, which submitted the UIA bid.
2. The 4 mini-VeriSigns, which run the new TLDs, including Register.Com (DotOrg and RegisterOrg), GNR, NeuStar, and Afilias (ISOC).
3. Everybody else. Unity was included in the table but the staff report went on at great length to discourage consideration of their bid.

Experience in the industry and size of the registry operated would be handy metrics to use, were it not for the fact that the average amount of experience of the 4 mini-VeriSigns is 7.8 months, the size of the databases they operate are small, and there is no public way of evaluating how well those systems have operated.

In fact, the rumors "on the street" are that the new TLDs have been beset by problems. Several requests for the voluminous reports on the operation of the new TLDs required in Appendix U[6] were met with protests from ICANN that the materials [might be confidential](#), needed to be reviewed by staff, and that staff didn't have the resources to do the review. We don't understand why standard reports on operations should have any confidentiality for a public utility. The reports required by ICANN do not serve the purpose of documenting and promoting transparency and there are no market pressures for VeriSign or the mini-VeriSigns to voluntarily do so.

A demonstrated ability to operate a registry service at an institutional level has not been shown by any of the applicants except VeriSign. Each of the other organizations bidding as operators are newly formed, somewhat undercapitalized startup ventures. We don't mean that as a criticism, only to point out that these are new organizations and the gTLDs they operate are new, small, growing slowly, and require a great deal of care and feeding from over-extended startup companies.

Staff reliance on simple uniform metrics submitted by the applicants with no verification of claims made a nice chart[5] (albeit a chart subject to [varying interpretations](#)), but this process didn't shed any insight into a "demonstrated ability" to operate a registry of "significant scale." By disqualifying IMS/ISC and bids such as Unity from serious consideration, the process has narrowed the choice to a small group of very similar bids.

3. Price

[TOC](#)

From May through September of 2002, the IMS/ISC team spent over 40 months of personnel time on the .org bid. While we suspect more of our team's time was spent on technical development than our competitors, it is clear that that many of our fellow supplicants spent a great deal of money as well. With 11 bidders, it is clear that the .org bidding process has consumed well over US\$5 million in resources, perhaps much more.

We respectfully submit that ICANN has succumbed to a form of .com excess. Tremendous resources have been expended in an excess of gold-rush opportunism, and the .org bid process is just one symptom of this excess. The DNS is a public service first, an industry second. While operating portions of the DNS can be moderately profitable if done correctly, the focus must be on long range (measured in decades, not months) plans, on conservatism instead of speculation. This is a public utility, not some go-go high-risk high-return market opportunity.

Inflated costs are reflected throughout the domain name process that ICANN manages. While we commend the ICANN board for their extremely low board expenses and the volunteer efforts of board members and we note that our experience with ICANN staff is that they work extremely hard, we do find it a bit incredible that an organization such as ICANN should, for the year ending June 30, 2001, spend US\$1,424,769 on a law firm, and US\$354,433 with Arthur Anderson.[7]

Hiring Gartner also reflects a .com philosophy. VeriSign and NeuStar are both large clients of Gartner. Like an investment bank, it is important that industry consultants properly disclose potential conflicts between their public analysis such as their analysis of the .org bids and their lucrative consulting engagements. Furthermore, the Gartner analysis displayed a shockingly naive approach to evaluating Internet infrastructure. A review by respected members of the Internet community willing to stake their reputation on their analysis would have yielded far superior results at a far lower cost and would have been far more accountable and transparent.

The high costs of the bidding process are a symptom of a larger problem. A wholesale cost of \$6 per domain name per year is far too high. Our bid was formulated knowing that MIS managers might be evaluating it, so we made sure to spec out customer support levels and hardware at a level that these people would understand as "stable." Even at these high cost levels, our consolidated financial analysis[8] shows that US\$3/name/year in 2003 with costs decreasing at 10% per year would more than cover the operating costs and a reasonable margin, even using old-line MIS techniques. With some careful provisioning, costs could easily be driven down to US\$2/name/year at the wholesale level. Some of our more radical colleagues suggested that these costs could be reduced by an order of magnitude more. We agree that some new gTLDs could easily cost less than \$0.25/name/year as a retail price to the consumer, but we also understand the conservatism needed to insure stability of the installed base of .org.

In our Supplemental Answers,[\[9\]](#) we stressed that one of our goals was the sale of fewer domain names and a desire to take the "MIS Mystique" out of the industry by producing standards compliant, high quality, open source software for registry operation (a process which is well underway despite the diversions of the .org bidding). The IMS/ISC pricing proposal was that we would charge no more than the lowest price on the market, and any surplus funds would get plowed into closely related infrastructure or returned to consumers. Indeed, our board has already been asking if we could simply price at US\$2/name/year in 2003. In short, we believe it is a worthwhile goal to help shrink the billion dollars per year DNS industry down to a smaller, more stable market size with a focus on service instead of speculative profits.

Each of the "qualified" bidders has a pricing philosophy that is very different from the one we proposed. [Dot Org Foundation](#), [GNR](#), [ISOC/Afilias](#), [Register.Org](#), [UIA/VeriSign](#), all proposed flat pricing of US\$6/name/year with many to-be-determined lucrative additional fee opportunities. [NeuStar](#) proposed a price of US\$5, but since they don't have to pay a commission to a nonprofit figurehead, their costs are lower.

The reason the wholesale price of a domain name hovers around US\$6/name/year is a reflection of ICANN's failure to create a competitive marketplace. Afilias is an alliance of registrars and includes a significant investment from VeriSign. Register.Com is, likewise, trying to make a significant business play in both registrar and registry portions of this market (though it should be noted that the financial markets [have not responded very enthusiastically](#) to this story).

Instead of a competitive marketplace, ICANN has created a small oligopoly. While this club is arguably open for new members to join, the application procedure requires such a large investment in time that it serves as a real barrier to entry. The lucrative opportunity to run a "real" registry is limited to those existing players, all of which are committed to keeping prices stable at US\$6/name/year. While an oligopoly of new small players might be considered preferable to the previous VeriSign monopoly, one has to wonder if there might not be yet other market structures if these barriers to entry were taken down.[\[10\]](#)

The effect of limiting the competition to the existing market players has a real cost to consumers. Our [consolidated financial analysis](#) projects a total of 3,846,000 billable events over the 5-year .org contract. If the bidding process results in a price of US\$6/name/year instead of US\$2/name/year the cost to consumers is US\$15,384,000.

The members of the current DNS oligopoly stand to lose far more total revenue if .org is operated by a new entrant who sets higher service levels with better reporting at lower prices than just the simple loss of the .org revenue stream. The .com/.net/.org installed base will have over 44,713,000 billable events over the next five years. Retail prices for domain names vary from US\$8.95 and up. Registrars are widely considered to be price insensitive (hence the remarkable stability of wholesale prices at US\$6), an indication that current margins are high and new players in the market would most surely drive prices down. If the artificial price boost from both the wholesale and retail markets amounts to US\$6, a fairly conservative number, the cost to consumers of these barriers to entry would be US\$268 million over the 5-year period.

4. Process

[TOC](#)

The .org TLD bid process generated an impressive amount of paperwork: Initial bid materials, proposals, presentations, a public forum, supplemental questions, a preliminary staff report with multiple evaluations, a second public forum, a final staff report with multiple revised evaluations, and a final response to the final report.

When we started preparing our bid, we decided that we would be extremely serious about following the bid procedures. Our team has extensive experience doing grants with ARPA, NSF, and foundations and has been involved in government contracting work for many years. We understand how RFPs work and we studied the materials quite carefully. We took great pains to understand the process and follow the rules carefully.

We were thus quite surprised to see the process unfold. Rather than allow supplicants to post their proposals on the web, ICANN insisted on repurposing all the content on their own web site, despite limited technical capacity to do so. On the .org applications web site, there are over 240 broken internal links caused by changing paths and a failure to properly unzip archives correctly, resulting in important material (such as all the supporting documentation for ISOC) to be inaccessible. No disrespect is intended to ICANN staff, who obviously work quite hard and are quite understaffed, it is simply a matter of attempting to build a custom experimental system where off-the-shelf techniques would be more suited to an operation of ICANN's scale.

Why not simply have each supplicant maintain their own web site and submit exactly one (potentially large) PDF file for the official application? ICANN staff seems quite sensitive to worries that supplicants will change or amend their bids, but the desire to freeze data at points in time, while a useful measurement technique, is not a substitute for a real evaluation.

Just as dialogue with the supplicants was carefully scripted through this long ever-changing process, dialogue with the community was also carefully managed. The [public forum for .org](#) was awkward and served mostly as a forum for PR

agents to publish press releases. Midway through the process ICANN staff shifted to a [second kind of public forum](#) that did not allow direct public posting. All data had to be sent via email for possible reposting, if deemed relevant, by ICANN staff. There are several very good open source bulletin board systems available that ICANN could install which would be a far more effective platform for supporting dialogue with the public.

Despite the difficulty of communicating with ICANN, a large community sprang up around the IMS/ISC bid. Over 600 people [voiced their support](#) and over 375 web sites [spread the dot](#). Looking at the web sites of the applicants didn't seem to be part of the evaluation process, but such a look sheds a lot of light about the bidders. The contrast between the true community effort that sprang up around the IMS/ISC bid and the contrived PR efforts of some of the other bids is quite evident.

These difficulties in communication seemed awkward, but surmountable. After all, our team has worked with extremely arcane communications software and systems from the earliest teletype-based word processors to the most hide-bound mainframe-based office "automation" systems. However, given the US\$29,000 examination fee, we really, really expected ICANN or its evaluators to do basic due diligence.

We mentioned the lack of technical due diligence in a letter to [our fellow .organisms](#) commenting on the draft staff evaluation. The response from ICANN staff to our protest was contained in the final staff evaluation:

Comment: No technical due diligence performed. Past performance not documented. No examination of codes, logs, or configurations.

Response: This was not within the scope of the RFP process. IMS/ISC did not comment to our knowledge that this should be part of the process when the draft RFP was posted for comment. IMS/ISC bid knowing that these factors were not part of the process. We cannot introduce into the evaluation requirements that were not part of the RFP.

Within available resources and timeframe, it is not likely that such a requirement could be performed. With regard to past performance, information was sought and obtained both in the application process and by supplemental questions. The necessity for, or even real value of, examination of codes, logs, or configurations is not clear. Moreover, such a highly intrusive exercise was not contemplated by the RFP and would risk constraining the range of reasonable engineering judgments.

On the question of financial due diligence, ICANN staff has a similar position:

No financial evaluations were made of any of the bidders. Demonstrated experience over extended periods was considered to be of greater significance since financial conditions can always change, positively or negatively.

As a reminder, the extended period of time in this case is under a year in the case of Afilias. In the case of the ISOC bid, the front-end provider is a so-called "Public Interest Registry" which does not yet exist and has not selected board members. Ignoring financial evaluations seems like a throwback to the height of the .com era when VCs would toss money at entrepreneurs based on a 3-slide presentation of a [novel business model](#).

(We should note in passing that a far more diligent evaluation was done by the EFF when we [asked them to support our effort](#). John Gilmore and Brad Templeton of the EFF fired dozens of questions at us about why we were bidding and exactly how we would spend the money, and protested that we should provision the service on [\\$2,000 Linux-based PCs](#). After a thorough evaluation of what we were proposing, they rejected our bid for support as being part of a process that was broken beyond repair.)

The lack of technical or financial due diligence is, quite simply, a disconnect. We don't understand it. It appears to us that ICANN has been highly focused on process over purpose, but has created a process so experimental and so heavyweight as to be unworkable. We realize that ICANN operates in a fishbowl, but the focus on producing a paper trail has led to huge expenditures of resources but shed very little light on the .org bids.

5. Parable

[TOC](#)

On June 14-15 1992, 11 people were [called into](#) a small conference room at the old National Science Foundation headquarters to evaluate proposals submitted to provide registry services for certain Top Level Domains.[\[11\]](#) The call for proposals was issued March 19, 1992 and bidders had two months to prepare their proposal.

Cooperative Agreement NCR-9218742 [\[12\]](#) was signed with Network Solutions, Inc. (NSI) and [announced](#) by the NSF when the system was operational on January 5, 1993 (albeit operational with [a few bugs](#)). Network Solutions was purchased by VeriSign March 7, 2000 [\[13\]](#) and is currently the largest player in the US\$1 billion/year "domain name industry."

The proposals received fell into two camps. Rick Adams, then of UUNET and now an IMS board member, submitted a [bare-bones simple proposal](#). We note that the Adams proposal called for a split between "registry" and "registrar" functions several years before this split was re-proposed in the White Paper.[\[14\]](#)

The other supplicants, including General Atomics, Network Solutions, AT&T, and one of the then-high-flying Big 5 firms all submitted full-blown proposals which waxed eloquent about enhanced information directory services and other buzzwords of that era. The review panel was intrigued by the Adams proposal, but ultimately threw it out because it was "too simple" and "could not be implemented with the costs proposed." None of the other proposals were terribly attractive, but the panel finally pulled a Solomon and split the award into three pieces, extracting one from each of three proposals.

General Atomics got to build an InfoScout, at the time deemed an essential element of the still-hazy comprehensive architecture for a "network information service." The InfoScout crashed and burned in the mid-term review.[\[15\]](#) AT&T was handed the piece for the Internic, a component that faded gracefully away to obscurity after having served no apparent purpose. Network Solutions got the boring job of running the database.

Reviewers were paid US\$100/day plus a modest expense stipend, though several reviewers didn't bother filing the paperwork for their remuneration. The bid process took a few months. No entry fee was required to submit a proposal. At least two reviewers wish to this day that they had pounded the table harder to have the Adams proposal considered much more seriously.

6. Purpose

[TOC](#)

Price, performance, and process are all means to an end. The .org reassignment is the culmination of a long path ICANN has traveled towards breaking up the inadvertent monopoly on registries granted to VeriSign. To put this event in perspective, here are some of the key events:

1. The registry function starts as a bookkeeping function. SRI gets saddled with the messy generic TLDs.
2. Network Solutions turns it into a service under NSF guidance.
3. The Internet starts to grow and the DNS turns out to be important.
4. A long somewhat murky road is followed to figure out what to do about the NSI monopoly and how to grow the DNS, which results in the creation of ICANN.
5. ICANN splits the market into wholesale and retail and starts accrediting registrars.
6. New TLDs are created with new registry operators. NSI, now called VeriSign, keeps running .com, .net, and .org as the registry and is still the largest registrar.
7. The .org reassignment attempts to break up the .com/.net/.org registry monopoly.

The reassignment of .org is a crucial milestone. However, .org is just one of several components of the top portions of the DNS. In addition to the new TLDs (.biz/.pro/.etc) and the big three (.com/.net/.org), there are a large number of ccTLDs, with which ICANN has a tenuous and somewhat nebulous relationship.[\[16\]](#)

One can argue that the creation of registrars or the new TLDs have been the key events shaping the market. However, the fact is that the bulk of the names are in the big 3. The new TLDs are on extremely shaky financial ground and will require a huge amount of care to grow, particularly with competition from the commercialization of ccTLDs such as .la and .us.

ICANN clearly wants to guide this market towards maturity, but it is doing so with a misguided sense of purpose. Limiting the .org reassignment to the mini-VeriSigns will not provide stability, and it will certainly not change how the market operates. Afilias, GNR, NeuStar, and Register.Com have participated heavily in ICANN, using the meeting ground provided by ICANN to become remarkably similar to VeriSign in their operational styles and pricing philosophies. VeriSign has been replaced by a small club, and membership in the club seems to have been the gating factor for qualifying to run .org.

There is a broader purpose and broader opportunity in reassigning .org than just assigning the associated revenue stream to one of five .coms. The .org TLD is a public trust. The RFP clearly stated a large number of goals, including stability, promoting competition, driving innovation, and serving the needs of the noncommercial world. The .org evaluations resulted in a simplistic evaluation that limited the choices to the mini-VeriSigns, with no consideration of these broader purposes.

At this point, it is useful to step back a bit and ask what ICANN's purpose is. We believe the functions of ICANN can be categorized as follows:

1. A role in the growth and regulation of TLDs.
2. A role in the coordination of protocol parameters.

3. A role in the coordination of numbers.
4. Overhead.

In the area of protocol parameters, the IETF has long been considered to be the custodian of the namespaces. The IETF has been particularly critical of ICANN's contribution in this area.[\[17\]](#) In the area of numbers, the heavy lifting has been done quite successfully by the well-established and well-respected number registries. The number registries have been particularly critical of ICANN's contribution in this area.[\[18\]](#) As previously noted, the ccTLDs have also been extremely critical of ICANN's role.[\[16\]](#)

That leaves overhead. We respectfully submit that ICANN is so consumed with process over purpose that the bulk of its resources have been devoted to overhead. Our attendance at the ICANN meeting in Bucharest and a close reading of ICANN-related materials such as the [ICANN public forums](#) and the General Assembly [mailing lists](#) confirms the impression that far more cycles are consumed on overhead than are devoted to useful work.

In short, ICANN has spent the bulk of its time in search of a purpose in life, "a business model" in .com-speak. Like the OSI standards effort, a focus on process over purpose has consumed resources but yielded few results. While we realize an intensive reform effort is underway, we respectfully submit that it appears to be simply more of the same, a view we believe we share with other organizations that interact with ICANN. We wouldn't be so presumptuous as to recommend that ICANN throw in the towel and disband, leaving others to sort out the mess, but it is clear that some form of reset is definitely in order.

One reset is to seriously consider the opportunity presented by the .org reassignment. Handing .org to a mini-VeriSign simply insures more of the same. High prices, low transparency, bad service, and lack of innovation are the result. A real evaluation of the bids shows several other viable options: ICANN can choose a go-go entrepreneur like Organic or a public service bid like IMS/ISC or a ccTLD expansion play like Unity. Each of these is a different model from the mini-VeriSign choices on the table.

A broader reset ICANN should definitely consider is a re-examination of how to grow the TLD space. While we disagree with those that advocate tens of millions of TLDs, we think there can be a solution that allows significantly faster growth but substitutes the current heavy-handed algorithm with something significantly more lightweight, based on interoperability requirements and points of presence instead of pitches and presentations. A neutral set of well-understood rules and clear points of interoperability will create more choice, better service, and lower prices for consumers. Given the all-consuming reform process currently underway, another organization, such as the IAB, might consider studying the issue and developing recommendations on a better algorithm for growth of the gTLD space that separates policy from process.

We close with a few words about our own purpose. Our [OpenReg software](#) is being readied for general release by January 1, 2003. Operation of the F root server, continued development of BIND, and a variety of other efforts are also coming along quite nicely. While it is a shame that consumers will be hurt by higher prices and little incentive to provide better service for gTLDs, there is plenty of work to do with ccTLDs, continued growth and evolution of the root servers, and better DNS service.

=====
ADVERTISEMENT
=====

The DNS is a public service. If you'd like to support the [Internet Software Consortium](#) through time, money, or partnerships, please contact Paul Vixie at vixie@isc.org.

=====
ADVERTISEMENT
=====

The IMS/ISC team has been quite gratified to see a [large community of supporters](#) form around our bid. We have always felt that any Internet effort is ultimately about the formation of true community. Operation of infrastructure in the public interest is part of a broad community effort. We greatly appreciate the support we've received from people throughout that broader community.

If you've been spreading the dot, you can [point your dot to the OpenReg pages](#) or, if you simply want to say a few choice words, feel free to [send mail to org-eval@icann.org](mailto:send_mail_to_org-eval@icann.org) to put it in the official ICANN queue. Thanks for the [dots!](#)

References

- [1] [Internet Multicasting Service](#) and [Internet Software Consortium](#), "[IMS/ISC .org Proposal](#)", June 2002.
- [2] [ICANN](#), "[Seventh Update on .org Reassignment Process](#)", September 2002.
- [3] [ICANN](#), "[Final Staff Evaluation Report on the Proposals for Reassignment of the .Org Registry](#)", September 2002.
- [4] [Internet Multicasting Service](#) and [Internet Software Consortium](#), "[The .org TLD is a Public Trust](#)", June 2002.
- [5] [ICANN](#), "[Supplemental Question #11 to .org Applicants](#)", September 2002.
- [6] [ICANN](#), "[Model Registry Agreement, Appendix U: Evaluation of Concepts](#)", June 2001.
- [7] [ICANN](#), "[IRS Form 990, Schedule A, Part II: Compensation of the Five Highest Paid Independent Contractors for Professional Services](#)", June 2001.
- [8] [Internet Multicasting Service](#) and [Internet Software Consortium](#), "[IMS/ISC Consolidated Pro Forma Financial Statements](#)", June 2002.
- [9] [Internet Multicasting Service](#) and [Internet Software Consortium](#), "[IMS/ISC Answers to Supplemental Questions](#)", September 2002.
- [10] [Froomkin, A.](#) and [M. Lemley](#), "[ICANN and Antitrust](#)", June 2002.
- [11] [US National Science Foundation](#), "[NSF Solicitation for Network Information Services Manager for NSFnet and the NREN](#)", NSF Solicitation 92-24, March 1992.
- [12] [US National Science Foundation](#), "[Network Information Services Manager\(s\) for NSFNET and the NREN: INTERNIC Registration Services](#)", NSF Cooperative Agreement NCR-9218742, March 1992.
- [13] [US Securities and Exchange Commission](#), "[Form 425: Filing of certain prospectuses and communications in connection with business combination transactions.](#)", Accession Number 0000950103-00-000362, March 2000.
- [14] [US Department of Commerce](#), "[Statement of Policy: Management of Internet Names and Addresses](#)", Docket Number 980212036-8146-02, June 1998.
- [15] [US National Science Foundation](#), "[Internic Mid-Term Evaluation and Recommendations](#)", December 1994.
- [16] [Council of European National Top-Level Domain Registries](#), "[Comments on ICANN Zone Access Policy](#)", September 2002.
- [17] [Internet Architecture Board](#), "[IAB Response to ICANN Evolution and Reform Committee Second Interim Report](#)", September 2002.
- [18] [RIPE](#), [APNIC](#) and [ARIN](#), "[Regional Internet Registries' Submission to the Committee on ICANN Evolution and Reform](#)", June 2002.

Authors' Addresses

[TOC](#)

Internet Multicasting Service
P.O. Box 217
Stewarts Point, CA 95450
US
Phone: +1.707.847.3720
Fax: +1.415.680.1556
URI: <http://not.invisible.net/>

Internet Software Consortium
950 Charter Street
Redwood City, CA 94063
US
Phone: +1.650.779.7000
Fax: +1.650.779.7055
URI: <http://www.isc.org/>

Appendix A. Word Count

[TOC](#)

```
./bin/xml2rfc full.stop.xml full.stop.txt ; \  
echo 'The number of words in this response is ' ; \  
perl -ne 'print if /Preamble$/ ... /Thanks for the dots/' full.stop.txt | \  
egrep -v "^Internet Multicasting.*\]" | wc -w ; \  
echo '.'
```

The number of words in this response is 5000



Appendix B. Excerpts from the Adams NIC Proposal, 5/1992

[TOC](#)

Executive Summary

We propose to implement a Master Information Registry to be accessed only by network service providers, rather than end users. This is a wholesale vs. retail approach. We expect end users to get BETTER response going through a "retailer" than under the existing arrangement. Additionally, we expect several organizations to offer registration services for a nominal fee to the public. This competition will increase service levels by providing a true alternative if registration service levels are unacceptable. Under this plan, it should be possible to get response time of one day or less.

We expect most, if not all existing network service providers (e.g. the NSF mid-level networks, the Commercial Providers, the International networks) to use this registration service directly.

This proposal provides for a single host acting as a "mutual exclusion" server for assigning unique names and numbers and acting as a master database. This model forces the distribution of the actual end user interaction to multiple servers working off of the single master copy. A distributed system is the only model that will scale long term.

The single host will be a fast workstation running a registration server (to be written as part of this proposal). Network service providers will interact with this server in real time, receiving an immediate response from the server rather than the existing seven days under the current model. The server will be intolerant of syntax errors or incomplete registrations. This greatly simplifies the implementation of the server. The network service provider is responsible for correctly filling out the registration template before submitting it for registration. This distributes the "human element" among many organizations and prevents a backlog at the actual registry. We will provide syntax checking software for Unix systems to be used by the network service providers. This client/server model also allows for the network service providers to create their own front ends that are more particularly suited to their needs.

An additional major feature of this approach is that it completely removes any policy matters from the actual registration process.

The transition from the existing DDN NIC is simple and straightforward. Under this plan, the existing NIC becomes the first client of the Information Registry. The two operations can be run in parallel until the new registration server is fully functional. At that point, the existing NIC cuts over to being a true client of the master registry and additional clients can be enabled.

Additionally, we will provide up to 1 gigabyte of disk space per year to be used as a master database. Again, no end user access is anticipated to this database. It should be distributed among the network service providers.

Project Requirement Plan

We will provide all of the registration services listed in the solicitation. Our model requires that we function as the Internet Registry and not as a delegate. Additionally, we will provide a master database server for the Desirable Database Services listed in section 2(C) of the solicitation.

We propose to implement a Master Information Registry ("MIR") to be accessed only by Network Service Providers ("Providers"). The MIR will not be accessible by end users. The end users will be required to register numbers and names through an authorized Provider.

This approach distributes the load among multiple Providers, with the MIR providing, a simple, small, fast mutual exclusion registry. We envision approximately twenty Providers around the world interacting with the MIR on behalf of their customers.

The MIR will be small and responsive. The MIR will be a single system allowing access only from qualified Providers in a strict format. No syntax errors or data omissions will be accepted. It is the responsibility of the Provider to verify syntax and completeness before submission. We will provide Unix software that can be used by the Providers before submission.

By allowing the MIR to be unfriendly towards errors, with no human intervention required, we can implement the server in a very straightforward manner. The response time for a correctly formed registration should be seconds. Additionally, by moving the human element from the server to the various Providers, most personnel costs can be eliminated.

We propose developing the software and operating the MIR for the first year at NSF expense. Future years will be at no cost to NSF and will be funded by annual fees from the Providers.

We will provide these services to all qualified Providers. We expect to do the qualification ourselves, but will also accept any specific recommendation from NSF or its designate.

This implementation totally separates the basic function of registration from any policy questions that may arise. It also eliminates the need for block allocation of numbers, as a network can be registered in "real time". This should reduce the number of numbers allocated, but unused.

This approach will reduce the time to register or update a name to seconds from the existing seven day (or MORE!) time frame.

Quality Monitoring

The MIR is basically a database system. A fundamental part of any database system is an activity log. This log is required for robustness and accountability. It also makes a wealth of information available for statistical analysis. We will produce reports showing availability, error rates, error types, response time and any other relevant information.

Appendix C. Document Formats

[TOC](#)

This document is available in the following formats:

- [\[html\]](#) This document.
- [\[xml\]](#) XML
- [\[txt\]](#) ASCII
- [\[nroff\]](#) NROFF
- [\[word\]](#) Word
- [\[pdf\]](#) PDF

[TOP](#)

[THE .ORG TLD IS A PUBLIC TRUST](#)

[»](#)