

IDNs: Straightforward Technical Problem or Machiavellian Nightmare?

Greg Goth



Three of the leading figures trying to solve the technical aspect of internationalized domain names (IDNs) – Internet domain names containing non-ASCII characters, such as those used in Arabic or Chinese – have been alternately hopeful and pessimistic recently. Vint Cerf, chairman of the ICANN board, says he's more optimistic about finally deploying a globally workable IDN solution than he's been in a year. Cary Karp, director of Internet strategy and technology at the Swedish Museum of Natural History in Stockholm, paints a darker picture of disingenuous and cynical maneuvering by parties with axes to grind. And John Klensin, former chairman of the Internet Architecture Board (IAB), says his outlook on one of the global Internet community's most vexing and longest-running problems depends on the developments on any given day.

"Things are moving fairly quickly right now in real time," Klensin says. "You could find days in which I would feel far more bitter, skeptical, and pessimistic than Cary, and you can find days on which Cary might sound fairly optimistic."

Within the past few weeks, both ICANN and the IAB have released working documents detailing the organizations' current IDN activities and recommendations. On 1 November, ICANN released its IDN "road map," detailing tests that have been concluded and those scheduled to run (www.icann.org/announcements/announcement-1-01nov06.htm). In announcing the road map, ICANN president and chief

executive officer (CEO) Paul Twomey said he expected the final tests and discussions around IDNs to be completed by the end of 2007.

Shortly before ICANN released its road map, the IAB issued a comprehensive evaluation, coauthored by Klensin, Karp, and Patrik Falstrom, examining the state of IDNs and the technology meant to enable them. Known as RFC 4690 (www.ietf.org/rfc/rfc4690.txt), the document details unforeseen issues that have arisen with established IDN standards and the possible next steps for deployment.

On the surface, therefore, you could surmise that the technical community was making steady progress in deploying new IDNs and discovering bugs as the global registries discovered confusions and incompatibilities in the panoply of languages, scripts, and intermediary coding schemes meant to ideally let anyone, anywhere, register a domain in their native tongue with as little confusion or modification as possible.

Stakes Go Up in Athens

However, IDNs' arcane technical requirements have long served as a wedge issue for governments, organizations, and individuals dissatisfied with the Internet's governance structure. These grumblings have become louder as the United Nations and its telecommunications regulatory arm, the International Telecommunications Union (ITU), have started heavily emphasizing the necessity of expanding the Internet's technology and gov-

ernance structure to all areas of the world. The ITU's main mechanism for asserting this has been meetings at the World Summit on the Information Society (WSIS) in Geneva 2003 and Tunis in 2005, and at the first Internet Governance Forum (IGF) held in Athens from 30 October – 2 November 2006 (www.intgovforum.org/).

The outgoing ITU secretary general, Yoshio Utsumi, wasn't regarded as a friend of ICANN or its procedures. In his opening remarks at the Athens meeting, he offered unvarnished criticism of the status quo and what could be interpreted as a challenge to the claims that the DNS hierarchy itself needed to be preserved:

"I do not share the perspective of those who argue that Internet governance is just a developing-country problem. I disagree, because the basis of this perspective [...] borders on arrogance. Many of the critics of the current system of Internet governance [...] understand that no matter what technical experts argue is the best system, or no matter what self-serving justifications are made that this is the only possible way to do things, there are no such systems or technologies that can eternally claim they are best."

However provocative Utsumi's remarks might have been, both Cerf and Andrzej Bartosiewicz, the rapporteur of the ITU's IDN working group (www.itu.int/ITU-T/studygroups/com17/idn/), say the new ITU hierarchy could be more conciliatory and collaborative in its approach to working with ICANN, the IETF, and regional and

IEEE Internet Computing

STAFF

Lead Editor: Rebecca L. Deuel
rdeuel@computer.org

Group Managing Editor: Steve Woods

Staff Editors: Kathy Clark-Fisher,
Jenny Stout, and Brandi Ortega

Production Editor: Monette Velasco

Magazine Assistant: Hazel Kosky
internet@computer.org

Graphic Artist: Alex Torres

Contributing Editors: Greg Goth,
Alison Skratt, and Joan Taylor

Business Development Manager:
Sandy Brown

Publisher: Angela Burgess,
aburgess@computer.org

Associate Publisher: Dick Price

**Membership/Circulation Marketing
Manager:** Georgann Carter

Advertising Supervisor: Marian Anderson

HOW TO REACH /C

Writers: Access www.computer.org/internet/author.htm. Articles are peer reviewed for technical merit and copy edited for clarity, style, and space. Unless otherwise stated, bylined articles and departments, as well as product and service descriptions, reflect the authors' or firms' opinion; inclusion in this publication does not necessarily constitute endorsement by the IEEE or the IEEE Computer Society.

Letters to the Editors: Send letters to lead editor Rebecca L. Deuel at rdeuel@computer.org.

On the Web: Access www.computer.org/internet/ or our online resource, *IEEE Distributed Systems Online*, at <http://dsonline.computer.org>.

Subscribe: Visit www.computer.org/subscribe/.

Subscription Change of Address: Send requests to address.change@ieee.org.

Missing or Damaged Copies: Contact help@computer.org.

Reprints of Articles: For price information or to order reprints, send email to internet@computer.org or fax +1 714 821 4010.

Reprint Permission: To obtain permission to reprint an article, contact William Hagen, IEEE Copyrights and Trademarks Manager, at copyrights@ieee.org.

Copyright and reprint permission: Copyright © 2007 by the Institute of Electrical and Electronics Engineers. All rights reserved. Abstracting is permitted with credit to the source. Libraries are permitted to photocopy beyond the limits of U.S. copyright law for patrons' private use those articles that carry a code at the bottom of the first page, provided the per-copy fee in the code is paid through the Copyright Clearance Center, 222 Rosewood Dr., Danvers, Mass. 01923.

governmental groups to quicken the pace of IDN deployment.

"If you read various resolutions and you read them sort of in a vacuum, for some of us it would appear the ITU continues to believe it has a self-given mandate to deal with a bunch of things that in the first order they rejected," Cerf says. "On the other hand, it also depends a great deal on personality, and things the new ITU secretary general [Hamadoun I. Toure] is saying sound a lot more cooperative and less competitive than his predecessor – and as a consequence, rather than reading the worst possible case into the resolutions, I'm hoping we see a more sensible way forward."

Bartosiewicz says the more significant change in the ITU hierarchy might be the new deputy secretary general for telecommunications, Houlin Zhao, who had been the director of the ITU's telecommunications standardization bureau.

"He was very interested in the Internet aspects of ITU activities," Bartosiewicz says. "From my perspective, he was a very good diplomat, not focused on fighting with ICANN, et cetera, but on cooperation."

Bartosiewicz says the ITU's likely role in the best possible circumstances would be one of advising member governments and Internet registries – especially those that, for whatever reason, have little communication with ICANN – on IDN deployment and compatibility issues; indeed, he says that registries comprised half the attendees at the latest ITU-IDN meeting.

However, Karp and Klensin both say that some of the loudest critics of the progress made thus far in deploying IDNs would use the ITU to fragment the DNS and shift control of domain administration from ICANN to the ITU. Regimes with an abiding interest in controlling the information flow into their nations could particularly threaten to split their portions of the Internet from the global network, using the ITU as a forum and the perceived slow progress on IDNs as a cudgel.

Although the entire DNS' interoperability remains one of the Internet community's philosophical touchstones, Klensin says regimes that are cynically criticizing internationalization efforts should be allowed to cut themselves off from the global network.

"As I said to Vint many times, I think countries have the absolute right to take themselves off the Internet if they so choose, and if a country is determined to do that, we should do whatever we can to educate them about the consequences," Klensin says. "And if they are still determined to do that, we should help them get off the Internet in ways that don't cause undue pain to everybody else. Then it becomes an economic and trade issue, and sooner or later they will come back or they won't."

Next Technical Steps

RFC 4690 outlines several problems that have been discovered since the specifications for Internationalized Domain Names in Applications (IDNA) – based on RFCs 3490, 3491, 3492, and 3454 – were published in 2003. Notably for end users, the fundamental approach of the IDNA 2003 standards allows almost all Unicode characters when converting local scripts to an IDN. In addition to unintended errors that might be caused by colliding characters, some might draw more sinister implications as RFC 4690 denotes:

It is now generally understood that, in addition to the collision problems of possibly equivalent words and hence labels, it is possible to utilize characters that look alike – 'confusable' characters – to spoof names in order to mislead or defraud users. That issue, driven by particular attacks such as those known as 'phishing,' has introduced stronger requirements for registry efforts to prevent problems than were previously generally recognized as important.

By turning around the approach that allows nearly all Unicode characters to instead employ a much more

restrictive set, and in which regional registries will be responsible for establishing and enforcing IDN registration policies, Cerf and Klensin believe the problem will be easier to solve. Klensin also says the publication of RFC 4690 will provide a clear rebuke to critics who claim that the IETF and ICANN are dragging their feet.

“Think of this as the tuning of a protocol that went out there a little bit too early, rather than a drastic turnaround,” Klensin says. “Given that you have a closed set of characters, at this point, if there’s doubt about a character getting in, you should leave it out. It’s protocol tuning, and I am optimistic enough about the general nature of the IETF process that I believe this problem would have been solved three years ago without so much pressure from a variety of actors to do something immediately. This time around there is a much broader understanding.”

“It is very frustrating to think about the potential hazards of confusability and the kinds of abuse consumers might be in for,” Cerf says, “and while you’re trying to inhibit that in the greatest way possible, to be beaten upon by some country representative somewhere complaining you are deliberately ignoring their interests or cultural interests – well, it’s an easy environment in which to stir friction. RFC 4690 and some of the exchanges that took place at the IGF went some distance toward allowing the impatient community to understand what some of the problems are, and the technical community to appreciate how passionate people are to get to some kind of a workable solution.”

Karp says that although he doesn’t discern any consensus at present, because whatever solution emerges will involve “a certain amount of pain for everybody,” the scope of the challenge is by no means insurmountable. “I don’t despair of human ingenuity being more than able to deal with this,” he says. ☐

Greg Goth is a freelance technology writer based in Connecticut.

News in Brief

The **W3C’s Extensible Stylesheet Language (XSL) working group** has updated its XML formatting standard to incorporate indexing, PDF bookmarks, revision bars, and other popular and previously proprietary extensions. It has also enhanced several existing **Extensible Stylesheet Language Formatting Objects (XSL-FO)** features, including offering more complex page numbering options and multiple color profile references. The XSL working group’s XSL-FO subgroup developed XSL-FO 1.1 with contributions from several industry representatives, including **Adobe Systems, HP, IBM, Isogen, PTC-Arbortext, and Sun Microsystems.**

More information is available at www.w3.org/TR/2006/REC-xsl11-20061205.

ICANN has approved the Asia-Pacific region’s first top-level domain registry. The **.asia domain**, which the **DotAsia Organization** will manage, brings the total number of Internet domain suffixes to 266.

More information is available at www.dotasia.org/about/intro.html.

Worldwide, nearly one-third of **imprisoned journalists** have published stories online, indicating that “authoritarian states are becoming more determined to control the Internet,” said Joel Simon, executive director of the **Committee to Protect Journalists**. The committee released its most recent analysis in December, noting that among the 24 nations that have imprisoned reporters, China topped the list with 31 jailed journalists, 19 of them Internet-based. The US has detained three journalists, including **Associated Press photographer Bilal Hussein**, who was taken into custody in Iraq nine months ago and has yet to be charged with a crime.

More information is available at www.cpj.org/Briefings/2006/imprisoned_06/imprisoned_06.html.

The US-based **Internet 2** unveiled the first major segment of its advanced net-

work at its annual member meeting in December. The segment has an initial **100-Gbit-per-second** capacity between New York, Washington, DC, and Chicago and is a first step in the project’s plan to increase networking capacity and flexibility for the research and education communities. Internet 2 also announced that **NYSERNet**, New York state’s research and education network consortium, has become the first regional network to connect to the Internet 2 Network infrastructure. Full network deployments is expected by mid-2007.

More information is available at www.internet2.edu and www.nysernet.org.

Following on its Washington, DC, launch in September, representatives of the **American National Standards Institute (ANSI), Standardization Administration of China (SAC), and the US National Institute of Standards and Technology (NIST)** gathered in Beijing in November to celebrate the official Chinese launch of the **Standards Portal**. The portal’s goal is to facilitate goods and services trade between the US and China. Among the portal’s features are dual-language information on the two nations’ standards systems and a database of 2,000 trade-related standards. It also offers access to 300,000 other national, regional, and international standards and guidelines.

More information is available at www.standardsportal.org.

According to a new report, **organized crime operations** are luring top **IT students** and graduates into the highly lucrative **Internet crime world** using KGB-like recruitment techniques from the Cold War era. The report, issued by **McAfee**, says that criminals are increasingly targeting universities and online forums in search of students and recent graduates to write viruses, launder money, and engage in identity theft.

More information is available at www.mcafee.com/us/about/press/corporate/2006/20061212_191010_e.html.