



IPv4 Address Sharing and Allocation for IPv6 Transition

Yong Cui • *Tsinghua University, China*

Wendong Wang and Qi Sun • *Beijing University of Posts and Telecommunications, China*

Lishan Li • *Tsinghua University, China*

Xingwei Wang • *Northeastern University, China*

Sharing IPv4 addresses is important for IPv4 service continuity when the Internet transitions to IPv6. The IETF proposes IPv4 address sharing during the transition period. This article discusses IPv4 address sharing principles, surveys IPv4 address sharing and allocation mechanisms, and presents IETF standardization efforts.

The problem of IPv4 address exhaustion has limited the development of the Internet for years.¹ IPv6, considered as the next-generation network layer protocol, isn't deployed widely enough to fully solve that problem. The smooth transition from IPv4 to IPv6 is one of the most critical issues for today's Internet. Operators and vendors have been investing in IPv6 transition, but the coexistence of IPv4 and IPv6 will last for a long time. As a result, operators have to guarantee IPv4 service continuity because a large number of applications are still IPv4 only, even though IPv4 addresses have run out.

To improve the use of the scarce network address during the IPv4 to IPv6 transition, the industry has extended the IPv4 address space by taking the transport layer port as an additional identifier. In this way, multiple subscribers can share one IPv4 address.² However, sharing an IPv4 address influences addressing, routing, and forwarding architectures. A challenging problem is to ensure that entities sharing IPv4 addresses experience all types of IPv4 services when transitioning to IPv6. It would be costly to have to adapt all the protocols, especially the newly designed IPv6 protocols, to the shared IPv4 address schema.

The IETF has been focusing on IPv4 address sharing since 2008. Various working groups, including Dynamic Host Configuration (DHC), Softwires, Port Control Protocol (PCP), and so on, are collaborating together to solve the problem.

The proposed solutions would affect not only the design and deployment of IPv6 transition technologies, but also the IPv4 address resource allocation scheme. The IETF community is trying to minimize the change to the current Internet architecture and develop the roadmap to the pure IPv6 world.

Here, we provide an up-to-date survey of the proposed IPv4 address sharing mechanisms in IETF. We analyze the design principles of the IPv4 address sharing mechanism, and abstract the characteristics when applying it during the IPv6 transition period. Considering the different IPv4 address sharing methods, the network architecture will be constructed in different schemes. According to this feature, we classify the address sharing solutions into two categories: carrier grade network address translator (NAT)-based, or carrier grade network (CGN)-based; and distributed port management. We then provide a comprehensive analysis to the solutions according to the classification. Last, we give an overview of the most recent actions in IETF related to IPv4 address sharing.

Design Principles of IPv4 Address Sharing

IPv4 address sharing enables subscribers to use the same IPv4 address with different transport layer ports, which might cause changes to the current IP architecture. To minimize effects on the traditional Internet architecture, a solution should consider the following critical design principles:

addressing, routing, forwarding, and subscriber awareness.

Addressing

Compared with traditional addressing schemes, IPv4 address sharing extends the address space by taking port information into consideration. Multiple subscribers can use the same public IPv4 address. The network identifier and locator are isolated in this case: The shared IPv4 address only plays the role of network identifier, while the underlay IPv6 address is the actual network locator. Subscribers sharing the same IPv4 address have to source all connections from the restricted ports unique to that address. Additionally, those devices shouldn't perform the duplicate address detection on the shared public IPv4 address.

Routing

Current routing protocols route traffic according to IPv4 addresses. In the IPv4 address sharing scheme, the packet should be identified with an IPv4 address and with transport layer port information, which causes incompatibility to the existing routing systems. To fully leverage the current practice, the dual-stack border router should support lightweight routing. Lightweight routing maps the IPv4 address plus port and IPv6 address on the border router, which transfers the IPv4 traffic to IPv6 and vice versa. The mapping bridges IPv4 and IPv6 routing systems in the context of IPv4 address sharing.

Forwarding

To cooperate with the lightweight routing, the dual-stack border router should employ Software technologies to forward the IPv4 traffic over IPv6 networks. On receipt of an IPv4 packet, the dual-stack border router checks the IPv4 address and port information against the established mapping table for the destination IPv6 address. The dual-stack border router then transforms the IPv4 forwarding to IPv6.

When receiving an IPv6 packet, the dual-stack border router can simply decapsulate or translate the packet to an IPv4 one for normal forwarding.

Subscriber Awareness

All of the proposed IPv4 address sharing solutions extend the IPv4 address space by introducing port information. When sharing IPv4 through CGN technology, the CGN manages the public IPv4 address and port in a centralized way and assigns the resource to an incoming session dynamically. The subscriber is not aware of the shared public IPv4 address when initiating a session. CGN can achieve a high sharing ratio, but suffers heavy logging and potential performance bottleneck. The end-to-end principle is also broken: sessions originated from the Internet can't reach the subscriber without NAT traverse technology. Alternatively, the operator can distribute the IPv4 address and available source ports to subscribers at the edge, which enables the subscribers to be aware of the external shared IPv4 address. The main advantage of this approach is that it offloads the centralized translation burden to end devices, and the subscriber keeps control over the public IPv4 address and port resource.

CGN-Based Address Sharing

The CGN-based solution introduces the CGN function on the network side, translating packets between private and public addresses at a per-flow level. Subscribers use private addresses for connections to the IPv4 Internet, while the CGN manages public IPv4 addresses and performs the NAT function. Because of the shortage of public IPv4 addresses, flows from multiple subscribers might share the same address with different ports.

Dual-Stack Lite is a typical CGN-based solution, which Figure 1 illustrates.³ Through combining CGN and IPv4-in-IPv6 tunneling, this solution achieves IPv4 address sharing

and promotes IPv6 deployment. The dual-stack border router, acting as a centralized CGN device, performs the network address and port translation (NAPT) function to share public IPv4 addresses. The CGN mapping table extends the NAPT translation table to include source IPv6 address, which distinguishes flows through a 5-tuple mapping table (source IPv6, source IPv4, source port, destination IPv4, and destination port). The dual-stack end devices (including end hosts or Customer-Premises Equipment at the edge) and border router are tunnel endpoints for IPv4 traffic traversing the IPv6 network. When traffic from end devices arrives at the border router, the border router decapsulates the packet to the original IPv4 packet and forwards it following the IPv4 route. For inbound traffic, the border router looks up the public IPv4 address and port in the extended NAPT table for the private IPv4 address, port, and IPv6 address. The border router tunnels the IPv4 packet to the IPv6 address specified in the matched entry.

The CGN-based solution assigns one port when a session request comes, achieving high sharing efficiency of the IPv4 address. The CGN function manages the IPv4 address resource in a centralized manner, which enables the operator to have full control over IPv4 address. However, the subscriber isn't aware of the external public IPv4 address and ports, which crashes some application types, including universal plug and play, peer-to-peer, and so on, and thus downgrades the user experience. Subscribers can leverage the protocols like PCP⁴ and Session Traversal Utilities for NAT⁵ to negotiate the external public IPv4 addresses and port(s), but with a high cost and complexity. In addition, the centrally-located border router has to dynamically maintain the per-flow level state. There could be a performance bottleneck because the border router might serve a large number of active session requests.⁶

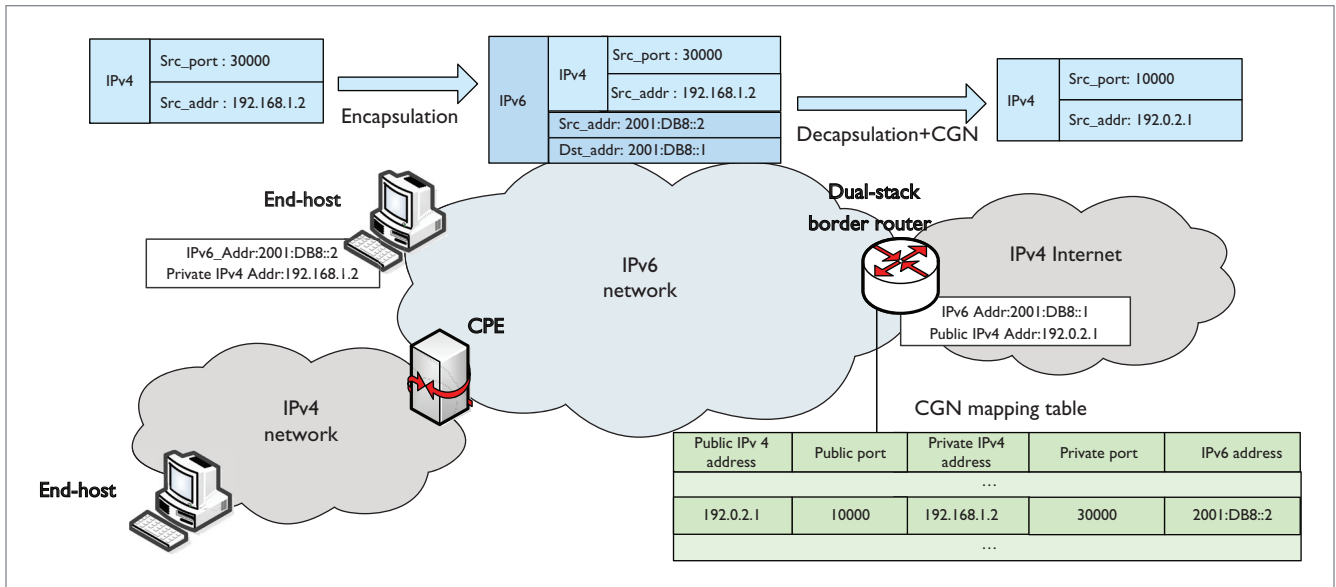


Figure 1. Carrier grade network address translator (NAT)-based, or carrier grade network (CGN)-based solution. The subscriber creates an IPv4-in-IPv6 tunnel ending at the border router when visiting the IPv4 Internet. The border router performs decapsulation/encapsulation and CGN functions, through the 5-tuple mapping table. The border router connects the IPv4 and IPv6 routing system.

Port Set Provisioning

The CGN-based solution can achieve a high sharing ratio of IPv4 addresses, but risking a potential performance bottleneck. The distributed port management solution offers optimization by moving the NAPT function from the centralized border router to the subscribers. This type of mechanism provisions public IPv4 addresses along with a nonoverlapped set of transport layer ports to a subscriber. Different subscribers use the same public IPv4 address but have source connections from different port sets. Subscribers have a relatively stable transport and IP layer configuration, which is essential to local NAPT functioning.

To provision IPv4 addresses and port sets to the end-user side, network entities must understand port sets. This requires an algorithm to represent a set of ports. On the other hand, it's critical to extend the current IP configuration protocol to support port set provisioning.

Algorithm for Port Set Representation

Port set is a set of contiguous/non-contiguous layer 4 source ports. The

IETF has proposed four algorithms for the calculation of available ports: Max-Min Range, Port Mask, Port Set Identifier (PSID), and Randomization.

The Max-Min Range algorithm specifies a continuous port range between the maximum and minimum port number.⁷ The Port Mask algorithm uses the mask operation to specify one or multiple port ranges.⁸ The port information is represented by two parameters: a port range value field that indicates the value of the significant bits of the port mask, and the port range mask that indicates the position of the significant bits for building the port mask. This algorithm is flexible but introduces unnecessary complexity. To preserve the system ports and simplify the mask algorithm, the PSID algorithm mandates contiguous mask bits like classless interdomain routing.⁹ This algorithm represents a port set with the following parameters: PSID offset, PSID length, and PSID value. The PSID offset is used to exclude the system ports. The PSID length specifies the length of the port set identifier. The PSID value specifies the value of the port set identifier assigned to the

subscriber. Considering attacks against transport protocols, the Randomization algorithm leverages the cryptographic mechanism to assign the random port numbers for each subscriber.⁸

The Max-Min Range algorithm is simple and straightforward, but it can only specify one continuous port range. Also, this algorithm might cause port range management issues and affect routing performance. The PSID algorithm is a subset of the Port Mask algorithm. But with this simplification, the PSID algorithm can easily exclude the system ports when the PSID offset field is larger than 0. Currently, the PSID algorithm is chosen by many distributed-port-management mechanisms, such as Lightweight 4over6,¹⁰ Mapping of Address and Port Using Translation (MAP-T),¹¹ and Mapping of Address and Port (MAP).⁹ The Randomization algorithm has high complexity and computing cost.

Predetermined Port Set Provisioning Protocol

The IPv4 address and port set (represented by PSID) can be provisioned in a predetermined or on-demand manner,

depending on requirements. The predetermined provisioning manner maps the IPv6 address/prefix to the IPv4 address and PSID before allocation, while the on-demand approach generates dynamic binding when requests come in.

The predetermined port set provisioning method has the advantage of simplicity and statelessness. IPv4 addresses and PSID parameters are explicitly bounded to the IPv6 prefix, or implicitly embedded in mapping rules for calculation. The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) is suitable for this provisioning pattern. The mapping information and rules are simply put into DHCPv6 option(s) and conveyed through DHCPv6 messages to the subscriber. DHCPv6 Options for Software is a typical predetermined port set provisioning protocol.¹² It defines DHCPv6 options that carry mapping rules, IPv4 to IPv6 bindings, and PSID parameters. The IPv4 address and PSID parameters can be retrieved following the MAP algorithm⁹ or explicitly configured.

DHCPv6 Options for Software mechanism configures the IPv4 address, port set, and mapped IPv6 prefix along with other DHCPv6 parameters, which leverages the DHCPv6 infrastructure. The features of predetermination and statelessness simplify the network management. But the utility of IPv4 address resources isn't efficient, requiring large IPv4 address space for the mapping. When the subscriber has exhausted the provisioned ports, the mechanism doesn't allow additional port set requests.

On-Demand Port Set Provisioning Protocol

On-demand port set provisioning protocols aim to make full use of the scarce IPv4 address resources. The IPv4 address as well as the port set is managed separately from the IPv6 prefix/address. The Dynamic Allocation of Shared IPv4 Addresses mechanism is designed for on-demand provisioning.¹³ This approach leverages DHCPv4 over

DHCPv6¹⁴ to dynamically manage the shared IPv4 address. The DHCPv4 part of the DHCP 4o6 server manages the PSID parameters with IPv4 address, and allocates PSID in a new DHCPv4 option. The provisioning process is similar to the dynamic allocation of full IPv4 addresses by a DHCPv4 server. DHCPv4 messages containing the IPv4 address and port set option traverse the IPv6 network using DHCP 4o6. Compared with the full IPv4 address allocation, the essential difference is that the same IPv4 address might be allocated to more than one subscriber.

This mechanism leverages DHCP 4o6 to dynamically provision the shared IPv4 addresses across IPv6 networks. IPv4 addresses and port sets can be allocated on-demand and the renewal/recycle procedure is independent from the underlay IPv6 address/prefix.

The PCP protocol can also support on-demand port set provisioning. The PCP port set mechanism extends the PCP protocol to allocate a port range, instead of a single port.¹⁵ This mechanism resolves the issue of source ports exhaustion: if additional ports are needed, the subscriber reinitiates a PCP request for another port set. For simplicity, the port set is represented through the Max-Min Range algorithm.

Applicability of Port Set Provisioning

Provisioning port sets ease the burden of IP address exhaustion. However, it's not universally applicable. This category of solutions should only be used on point-to-point links/tunnels. It's not suitable for network access over shared media, including Ethernet, WLAN, cable, and so on.

The port set provisioning mechanisms are typically used with the A+P IPv6 transition technologies developed in the IETF Software working group. MAP, MAP-T, and Lightweight 4over6 offload the centralized NAT by provisioning shared IPv4 addresses. MAP and MAP-T adopt mathematic mapping between IPv4 and IPv6, achieving

stateless maintenance at the operator side. Both use the predetermined port set provisioning manner, that is, Software DHCPv6 options. When there is no IPv4 embedded in the IPv6 prefix, DHCP 4o6 is also applicable for MAP. Lightweight 4over6 advocates simple one-to-one binding between the IPv4 information and IPv6 address without complex calculation. It's compatible with both the predetermined and the on-demand port set provision mechanisms. It can be deployed with any protocol of DHCPv6, DHCP 4o6, or PCP.

Figure 2 shows how to provision port sets in the scenario of IPv4 over IPv6. The provisioning system allocates the IPv4 address and port set over a IPv6 network in a stateless or stateful manner. The border router (BR) installs mathematic MAP rules for MAP domains if it works as a MAP BR, or it establishes the v4 to v6 mapping of the subscriber's public IPv4 address, port set, and source IPv6 address, as in Lightweight 4over6. With the rules and mappings as the heterogeneous routes, the BR connects IPv4 islands through IPv4-over-IPv6 Softwires.

IETF Standardization

Because of IPv4 exhaustion, IPv6 transition protocols have to consider IPv4 address sharing. Several IETF working groups, including Software, DHC, PCP, and Intarea, have participated in the development of IPv4 address sharing and the related IPv6 transition protocols.

Software focuses on the architecture development of IPv6 transition technologies. The A+P architecture sketches the scheme of IPv4 address sharing in the context of IPv6 transition.² To complete the architecture with signaling element, the Software Working Group is publishing *Software DHCP Options*. The DHC and PCP working groups extend DHCP and PCP protocols for allocating port sets, respectively. For example, the DHC Working Group has approved the mechanism of dynamic allocation of a shared IPv4 address. Additionally, the Intarea

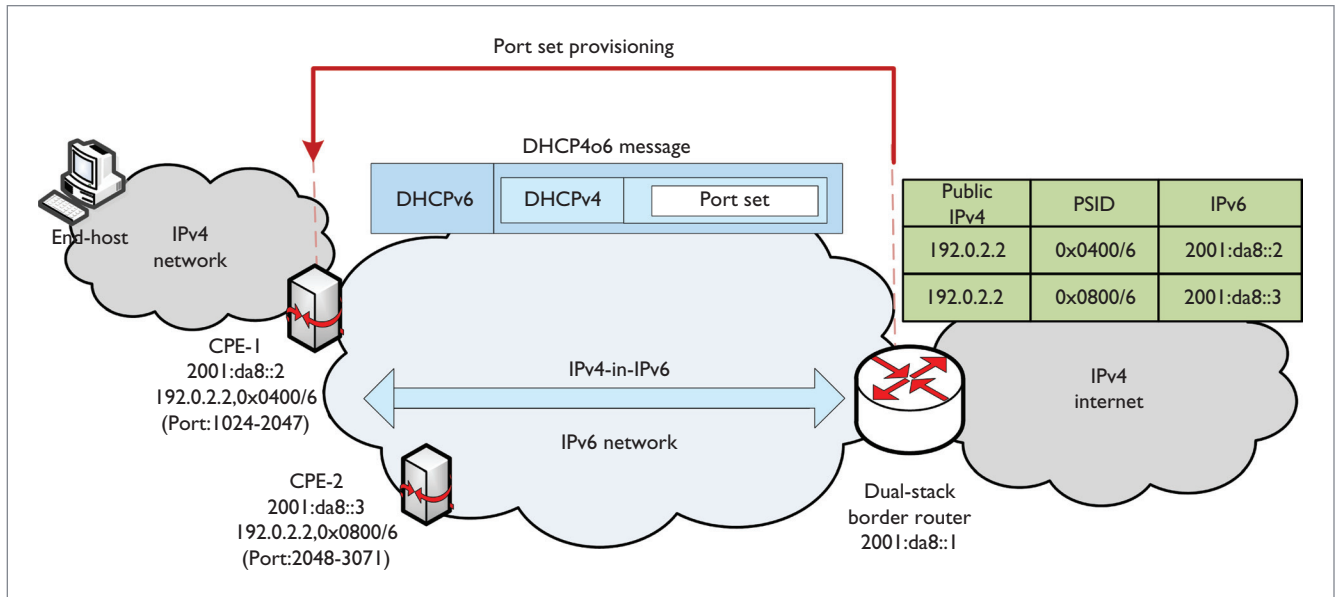


Figure 2. Port set provisioning using the Dynamic Host Configuration Protocol (DHCP) 4o6. The port set parameters are encoded as a DHCPv4 option and provisioned with the IPv4 address over IPv6. The port set is calculated using the Port Set Identifier algorithm (0x8000/6 means the first 6 bits are PSID). The border router is configured with corresponding v4 to v6 mappings for encapsulation/decapsulation.

Table 1. Comparison of IPv4 address sharing solutions.

Solution	Subscriber awareness of external v4	Port resource pattern	Allocation	State maintenance	Sharing efficiency	Logging	Provisioning method*
Dual-Stack Lite	Unaware	Port number	On-demand	Per session	High	Heavy	N/A
Lightweight 4over6	Aware	Port set	Predetermined/on-demand	Per subscriber	Low	Light	DHCPv6-based PCP-based DHCP4o6-based
Mapping of address and port	Aware	Port set	Predetermined	Stateless	Low	Light	DHCPv6-based
Mapping of address and port using translation	Aware	Port set	Predetermined	Stateless	Low	Light	DHCPv6-based

* DHCP = Dynamic Host Configuration Protocol; PCP = Port Control Protocol.

Working Group analyzes the potential issues of IPv4 address sharing.¹⁶ Table 1 compares the major IPv6 transition protocols with IPv4 address sharing. Operators must consider the available IPv4 address space, the IPv4 configuration protocol and the IPv6 transition network structure when choosing a port set allocation mechanism.

Sharing IPv4 addresses is an important part for IPv6 transition. CGN

is an innovative trial for enabling sharing IPv4 management in the IPv6 transition, but with apparent flaws. Port set provisioning is now the focus of IETF when developing IPv6 transition technologies. Currently, the IETF v6ops Working Group is going to summarize the deployment and operational experience for the IPv6 transition protocols with IPv4 address sharing. The goal is to transition to the pure IPv6 network as soon as possible. □

Acknowledgment

The National Natural Science Foundation of China (grants 61422206, 61120106008, 61225012) supported this work.

References

1. P. Wu et al., "Transition from IPv4 to IPv6: A State-of-the-Art Survey," *IEEE Comm. Surveys & Tutorials*, vol. 15, no. 3, 2013, pp. 1407–1424.
2. R. Bush, ed., *The Address Plus Port (A+P) Approach to the IPv4 Address Shortage*, IETF RFC 6346, Aug. 2011.

3. A. Durand et al., *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*, IETF RFC 6333, Aug. 2011.
4. D. Wing et al., ed., *Port Control Protocol (PCP)*, IETF RFC 6887, Apr. 2013.
5. J. Rosenberg et al., *Session Traversal Utilities for NAT (STUN)*, IETF RFC 5389, Oct. 2008.
6. S. Perreault et al., *Common Requirements for Carrier-Grade NATs (CGNs)*, IETF BCP 127, IETF RFC 6888, Apr. 2013.
7. T. Lemon, Y. Lee, and P. Wu, *Dynamic Host Configuration Protocol (DHCP) Options for Port Set Assignment*, IETF Internet draft, work in progress, Apr. 2012.
8. G. Bajko et al., *Port Restricted IP Address Assignment*, IETF Internet draft, work in progress, Apr. 2012.
9. O. Troan and W. Dec, *Mapping of Address and Port with Encapsulation (MAP-E)*, IETF RFC7597, July 2015.
10. Y. Cui et al., *Lightweight 4over6: An Extension to the DS-Lite Architecture*, IETF RFC7596, July 2015.
11. X. Li and C. Bao, *Mapping of Address and Port Using Translation (MAP-T)*, IETF RFC7599, July 2015.
12. T. Mrugalski, *DHCPv6 Options for Configuration of Software Address and Port Mapped Clients*, IETF RFC7598, July 2015.
13. Y. Cui and Q. Sun, "Dynamic Allocation of Shared IPv4 Addresses," IETF Internet draft, work in progress, May 2015.
14. Q. Sun et al., *DHCPv4-over-DHCPv6 (DHCP 4o6) Transport*, IETF RFC 7341, Aug. 2014.
15. Q. Sun and M. Boucadair, *Port Control Protocol (PCP) Extension for Port Set Allocation*, IETF Internet draft, work in progress, May 2015.
16. M. Ford et al., ed., *Issues with IP Address Sharing*, IETF RFC 6269, June 2011.

Yong Cui is a full professor at Tsinghua University, China. His research interests include computer network architecture and mobile computing. Cui has a PhD in computer science from Tsinghua University. Contact him at cuiyong@tsinghua.edu.cn.

Wendong Wang is a full professor at Beijing University of Posts and Telecommunications. His research interests include future Internet architecture, software-defined networking


technologies, and network or service quality of service. Wang has an ME in computer science from Beijing University of Posts and Telecommunications. He's contributed to three IETF drafts related to the Dynamic Host Configuration Protocol. Contact him at wdwang@bupt.edu.cn.

Qi Sun is a master's student at Beijing University of Posts and Telecommunications. His research interests include IPv4 to IPv6 transition and DHCP. Sun has a BS in communication engineering from Beijing University of Posts and Telecommunications. He has published two IETF RFCs on IPv6 transition technologies. Contact him at sunqibupt@gmail.com.

Lishan Li is a PhD candidate at Tsinghua University, China. Her research interests

include IPv4 to IPv6 transition, DHCP, and 464XLAT. Li has a BS in computer science from the University of Science and Technology Beijing. She has contributed one IETF draft related to the 464XLAT protocol. Contact her at lilishan48@126.com.

Xingwei Wang is a full professor at Northeastern University. His research interests include future Internet, cloud computing, and information security. Wang has a PhD in computer science from Northeastern University. Contact him at wangxw@mail.neu.edu.cn.

 Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.



Call for Articles

IEEE Software seeks practical, readable articles that will appeal to experts and nonexperts alike. The magazine aims to deliver reliable information to software developers and managers to help them stay on top of rapid technology change. Submissions must be original and no more than 4,700 words, including 200 words for each table and figure.

Author guidelines:
www.computer.org/software/author.htm
Further details: software@computer.org
www.computer.org/software

IEEE Software