IPv6 - Good for Grid: a position statement based on a technical viewpoint

S. N. Bhatti *P. Kirstein* *P. Clarke* *T. Chown* *D. Hutchison*
S.Bhatti@cs.ucl.ac.uk P.Kirstein@cs.ucl.ac.uk clarke@hep.ucl.ac.uk tjc@ecs.soton.ac.uk dh@comp.lancs.ac.uk

NETSYS http://www.netsys.ucl.ac.uk/
University College London
Gower Street
London WC1E 6BT, UK

ECS
Southampton University
Southampton
SO17 1BJ, UK

Computing Dept.
Engineering Building
Lancaster University
Lancaster, LA1 4YR
UK

Keywords: IPv6, Grid, routing/addressing, configuration, quality of service (QoS)

IPv6 offers great benefits to the Grid infrastructure and Grid applications, and so the Grid community should support research in this area. This is discussed briefly in the context of several key features of IP networking, namely: ***addressing and routing, mobility, end-system configuration, high-performance communication, quality of service (QoS) support, provision of security services, group communications and transition***.

Lack of IPv4 address space and scalability of routing

Ideally, every IP connected device should have a unique IP address. This is currently not practical with the IPv4 32-bit address space. In order to try and solve the problems arising from lack of addresses, mechanisms like Network Address Translation (NAT) coupled with Application Level Gateways (ALGs), such as web-proxies, are used. However, while such mechanisms are useful, they disrupt the end-to-end service semantics of application-level communication; which has implications for services such as security, mobility and QoS.

IPv6 has a 128-bit address space and uses address aggregation extremely well. For example, for IPv4, UKERNA currently has to advertise hundreds of network addresses in its routing tables for UK sites. With IPv6, UKERNA has a single address space - 2001:0630::/32 – for all UK sites. This improves routing information exchange and the operation of routing algorithms. There is also discussion of how to provide multi-homing and mobility in a transparent manner for IPv6 nodes by using different addressing schemes with IPv6 addresses.

In an IPv6 Grid, every node, whether database server, supercomputer or handheld device, could have a globally unique IPv6 address, irrespective of where it is located. A Grid infrastructure based on IPv6 allows direct IP-level connectivity of more end-systems and better routing scalability compared to IPv4.

Mobility support for end-systems at the IP level

In IPv6, there is explicit support for mobile end-systems and applications. There are aspects of the mobile support that are still being refined, but certainly it will be possible for mobile IPv6 systems to have globally unique IP addresses, benefit from optimised routing for mobile communication paths in the IPv6 domain and make use of group communication when using mobile connectivity.

In an IPv6 Grid, every pair of nodes, whether each is database server, supercomputer or handheld access device, can have a unique relationship – irrespective of how each is attached to the Internet. Indeed, because of the availability of group addresses, this relationship can be defined for whole groups of entities.

Configuration of end-systems at the IP level

With IPv4, there is no standard mechanism for an end-system to be configured automatically with basic information such as an IP-address, IP net-mask and default router. Also, automatic and transparent reconfiguration of the end-system is not easily supported in IPv4. An IPv4 network must rely on a knowledgeable administrator using additional protocols (e.g. DHCP) to set-up configuration services.

With IPv6, sensible auto-configuration options (including link-local and site-local addressing) are supported directly by end-systems and IPv6 routers. This allows an IPv6 end-system to be "plug-and-play" with respect to IPv6 connectivity. For example, a cluster of machines could all be plugged in and ready for communication in minutes. IPv6 auto-configuration can work with or without the presence of an IPv6 edge-router. DHCPv6 exists to extend this basic IP-level configuration capability to a rich set of configuration services.

Easy IP-level configuration capability would allow Grid users to construct and use more easily an IP-based infrastructure without having to have detailed network-specific knowledge.

Performance factors

Mechanisms implemented in the network layer in IPv4 have been streamlined and improved for IPv6. These modifications potentially allow much faster packet-handling of IPv6 packets compared to IPv4. The IPv6 header is much simplified compared to IPv4. This makes it more amenable to hardware processing. Also, IP-layer header checksums, fragmentation and reassembly have been removed from IPv6 routers. Header processing has been much improved by using chained headers with headers ordered sensibly to allow frequently processed headers to appear nearer the start of the packet while end-to-end headers appear just before the IPv6 payload. Detractors of IPv6 have stated that the performance of routers has been much poorer throughput with IPv6 than IPv4. This performance hit has been due to the lack of hardware support for handling the larger IPv6 addresses; recent advances in such support allow IPv6 packets to be handled at line speeds.

The simplified packet format, router functions and header processing for normal packet handling in IPv6 should allow improvements in performance for Grid applications.

Support for QoS sensitive flows

IPv4 was never designed to support real-time, interactive, QoS-sensitive applications. However, applications sensitive to delay, loss and changes in throughput will run on the Grid. There has been work within the IETF over many years to try and define QoS support and resource control in the network and in IP. Efforts such as INTSERV and DIFFSERV are all aiming to improve the QoS experienced by IP flows. Indeed, INTSERV and DIFFSERV are applicable to both IPv4 and IPv6; so it could be argued that there will be no advantage to IPv6 in terms of QoS support. However, coupled with the simplified header and better addressing capability, we should see improved performance with IPv6, and the flow-label in the header has the potential to allow additional QoS support. For example, by uniquely identifying end-to-end flows in the IPv6 header, rather than using transport- or link-layer header fields, we could simplify the processing of QoS sensitive flows in a router.

Some Grid applications would like to have protected flows in the network. Such QoS-sensitive Grid applications could benefit from using the IPv6 flow label to identify individual end-to-end flows.

Support for security services

With Grid users potentially being widely geographically dispersed, and the use of dynamically formed virtual communities, being able to exchange information securely across the public Internet will be of great importance. So, security infrastructure is required. Much of this security infrastructure will be at the application-level including public-key infrastructure and security management capability. Also, many security services are likely to be application-specific. However, there will be requirements for basic, common security services to allow packets to be transmitted securely across the network. So, the IETF IPSEC WG has defined mechanisms to allow IP packets to be shipped with authentication headers (IPSEC-AH) as well as protected against inspection (IPSEC-ESP). These mechanisms are mandatory in full IPv6 implementations, but optional in IPv4. The mechanisms to be used in mobile/wireless services have not been fully agreed yet. Nevertheless, there is potential for true end-to-end security services directly to IPv6 end-systems that are likely to be unavailable in the IPv4 regime in the same manner.

Grid applications using IPv6 will be able to use a common, base set of IP-level security mechanisms.

Support for group services

The IPv6 specifications allow a much cleaner definition of multicast addresses, and a new *anycast* address. These will allow new services to be defined at the network level, both for resource discovery and for group operations. It should be possible even to define private networks based only on such group addressing – irrespective of individual IP entities. The range of such services has not yet been defined fully, but is likely to have immense benefit to the Grid community, especially for applications such as dynamic software updates or multipoint-to-multipoint data delivery for distributed applications.

Consistent use of IPv6 multicast and anycast will allow powerful new mechanisms for the Grid community

IPv6 deployment issues

There are still some technical problems to resolve with IPv6, but these are being addressed and there is a world-wide community of IPv6 users – see http://www.6bone.net/ and http://www.6net.org/ . Equipment vendors can make money by building products and services that extend the life of IPv4. However, these engineering solutions are all retrospective additions and extensions to a protocol that is becoming increasingly more complex to use and manage. Customer demand for IPv6 is going to be one way of shifting the *status quo*: vendors do not produce IPv6 equipment for lack of customer demand, and customers do not demand IPv6 as they are offered interim solutions that use IPv4. IPv4/IPv6 transition methods have been defined and are being deployed.

While IPv4 will remain with us for a long time, deployment of transition aids from IPv6 will ensure that legacy applications and networks will continue to be accessible – albeit with inferior functionality.

History suggests that, to break the chicken-and-egg situation, aggressive adoption and deployment of IPv6 will encourage development of IPv6 applications. There will then be no shortage of people, from demanding e-Science users to commercial developers, who will find ways of exploiting the functionality to its limit.

Vitae

Dr. Saleem N. Bhatti is Lecturer in the Department of Computer Science and is the coordinator of the UCL NETSYS Group. He is involved in several Grid/e-Science projects at UCL, specifically in the area of networked systems. His work has covered a wide range of networked systems topics, including multi-service networking, tele-working, multicast, network and systems management, network protocol design, network and application security, IPv6 and QoS (Quality of Service) adaptability support for Internet applications and adaptive systems.

Prof. Peter Kirstein is Professor of Computer Systems and Director of Research in the Department of Computer Science, UCL. He has been involved in research and development of large IP-based networks, IP-based protocols and applications for over 30 years, introducing the first Internet node to the UK in 1973. He is the UCL contact for the 6NET project, which will provide extremely high-speed (Gb/s) native IPv6 capability on an international scale. He is proposed as the Director of the UK IPv6 Forum, which is one of the National IPv6 fora being set up.

Prof. Peter Clarke Professor of Physics at UCL and is involved in several Grid/e-Science project and initiatives. He is also Chair of the PPNCG (Particle Physics Network Coordinating Group) and Deputy Chair of the UK Grid Network Team (GNT). He is a member of the Global Grid Forum Governing Body (the Grid Forum Steering Group) and Director of the Data Area for the GGF. Prof. Clarke is also WP7 Deputy Chair (Networking) in the EU-DataGRID project and involved with networking activities in the DataTAG project.

Dr. Tim Chown is a lecturer, researcher and systems manager at the Department of Electronics and Computer Science at the University of Southampton. He leads the Department's participation in a number of projects focusing on IPv6, including 6INIT, 6WINIT and the new flagship EU IST projects 6NET and Euro6IX. Dr. Chown also leads UK academic IPv6 pilot activities and the IPv6 Working Group of the European GEANT project. He acts as an advisor to UKERNA in areas including IPv6, network QoS and H.323 conferencing, and is a member of the UK Grid Network Team. He also has interests in areas including wireless networking, security, messaging systems, IP multicast and multimedia applications.

Prof. David Hutchison David Hutchison is Professor of Computing at Lancaster University. His research interests are architecture, services and protocols for distributed multimedia systems, including QoS for the Internet, and he has been involved in many UK and European collaborative projects in these areas. He is a programme committee member for several key international conferences and workshops, including IEEE INFOCOM and ACM SIGCOMM; he was Programme Chair of IEEE OpenArch 2000 and Programme Co-Chair of IWQoS 2001. As of mid 2001, he has been a member of the UK Grid Technical Advisory Group (TAG), and Chair of the UK Grid Network Team (GNT).