

Deployment of 464XLAT (RFC6877) alongside IPv6-only CPU resources at WLCG sites

Terry Froy, Daniel Traynor, Christopher J. Walker
School of Physics and Astronomy, Queen Mary University of London

We utilized cutting-edge transition technologies to provide IPv6-only WLCG computing resources with full backwards-compatible connectivity to IPv4-only WLCG resources.

Motivation

- 1) IPv6 (Internet Protocol version 6) has been production-ready for more than a decade but not all WLCG sites are offering services over IPv6 yet.

3) A desire to measure how many IPv4-only hosts are still operated by the WLCG that are critical to the correct execution of production workloads; as far as we are aware, QMUL is the only WLCG site operating IPv6-only WNs on production WLCG queues.
- 2) Mandated requirement for all WLCG sites to be able to support the use of IPv6-only Worker Nodes (WNs) by April 2017 and the Tier-2 site based at QMUL wanted to adopt and support this strategy.

4) There is no definitive list of IPv4-only hosts which are required to have IPv6 connectivity in order for IPv6-only WNs to be used on the WLCG without any transition technologies such as 464XLAT; the logs from our NAT64 gateway cluster are intended to provide such a list.

Solution

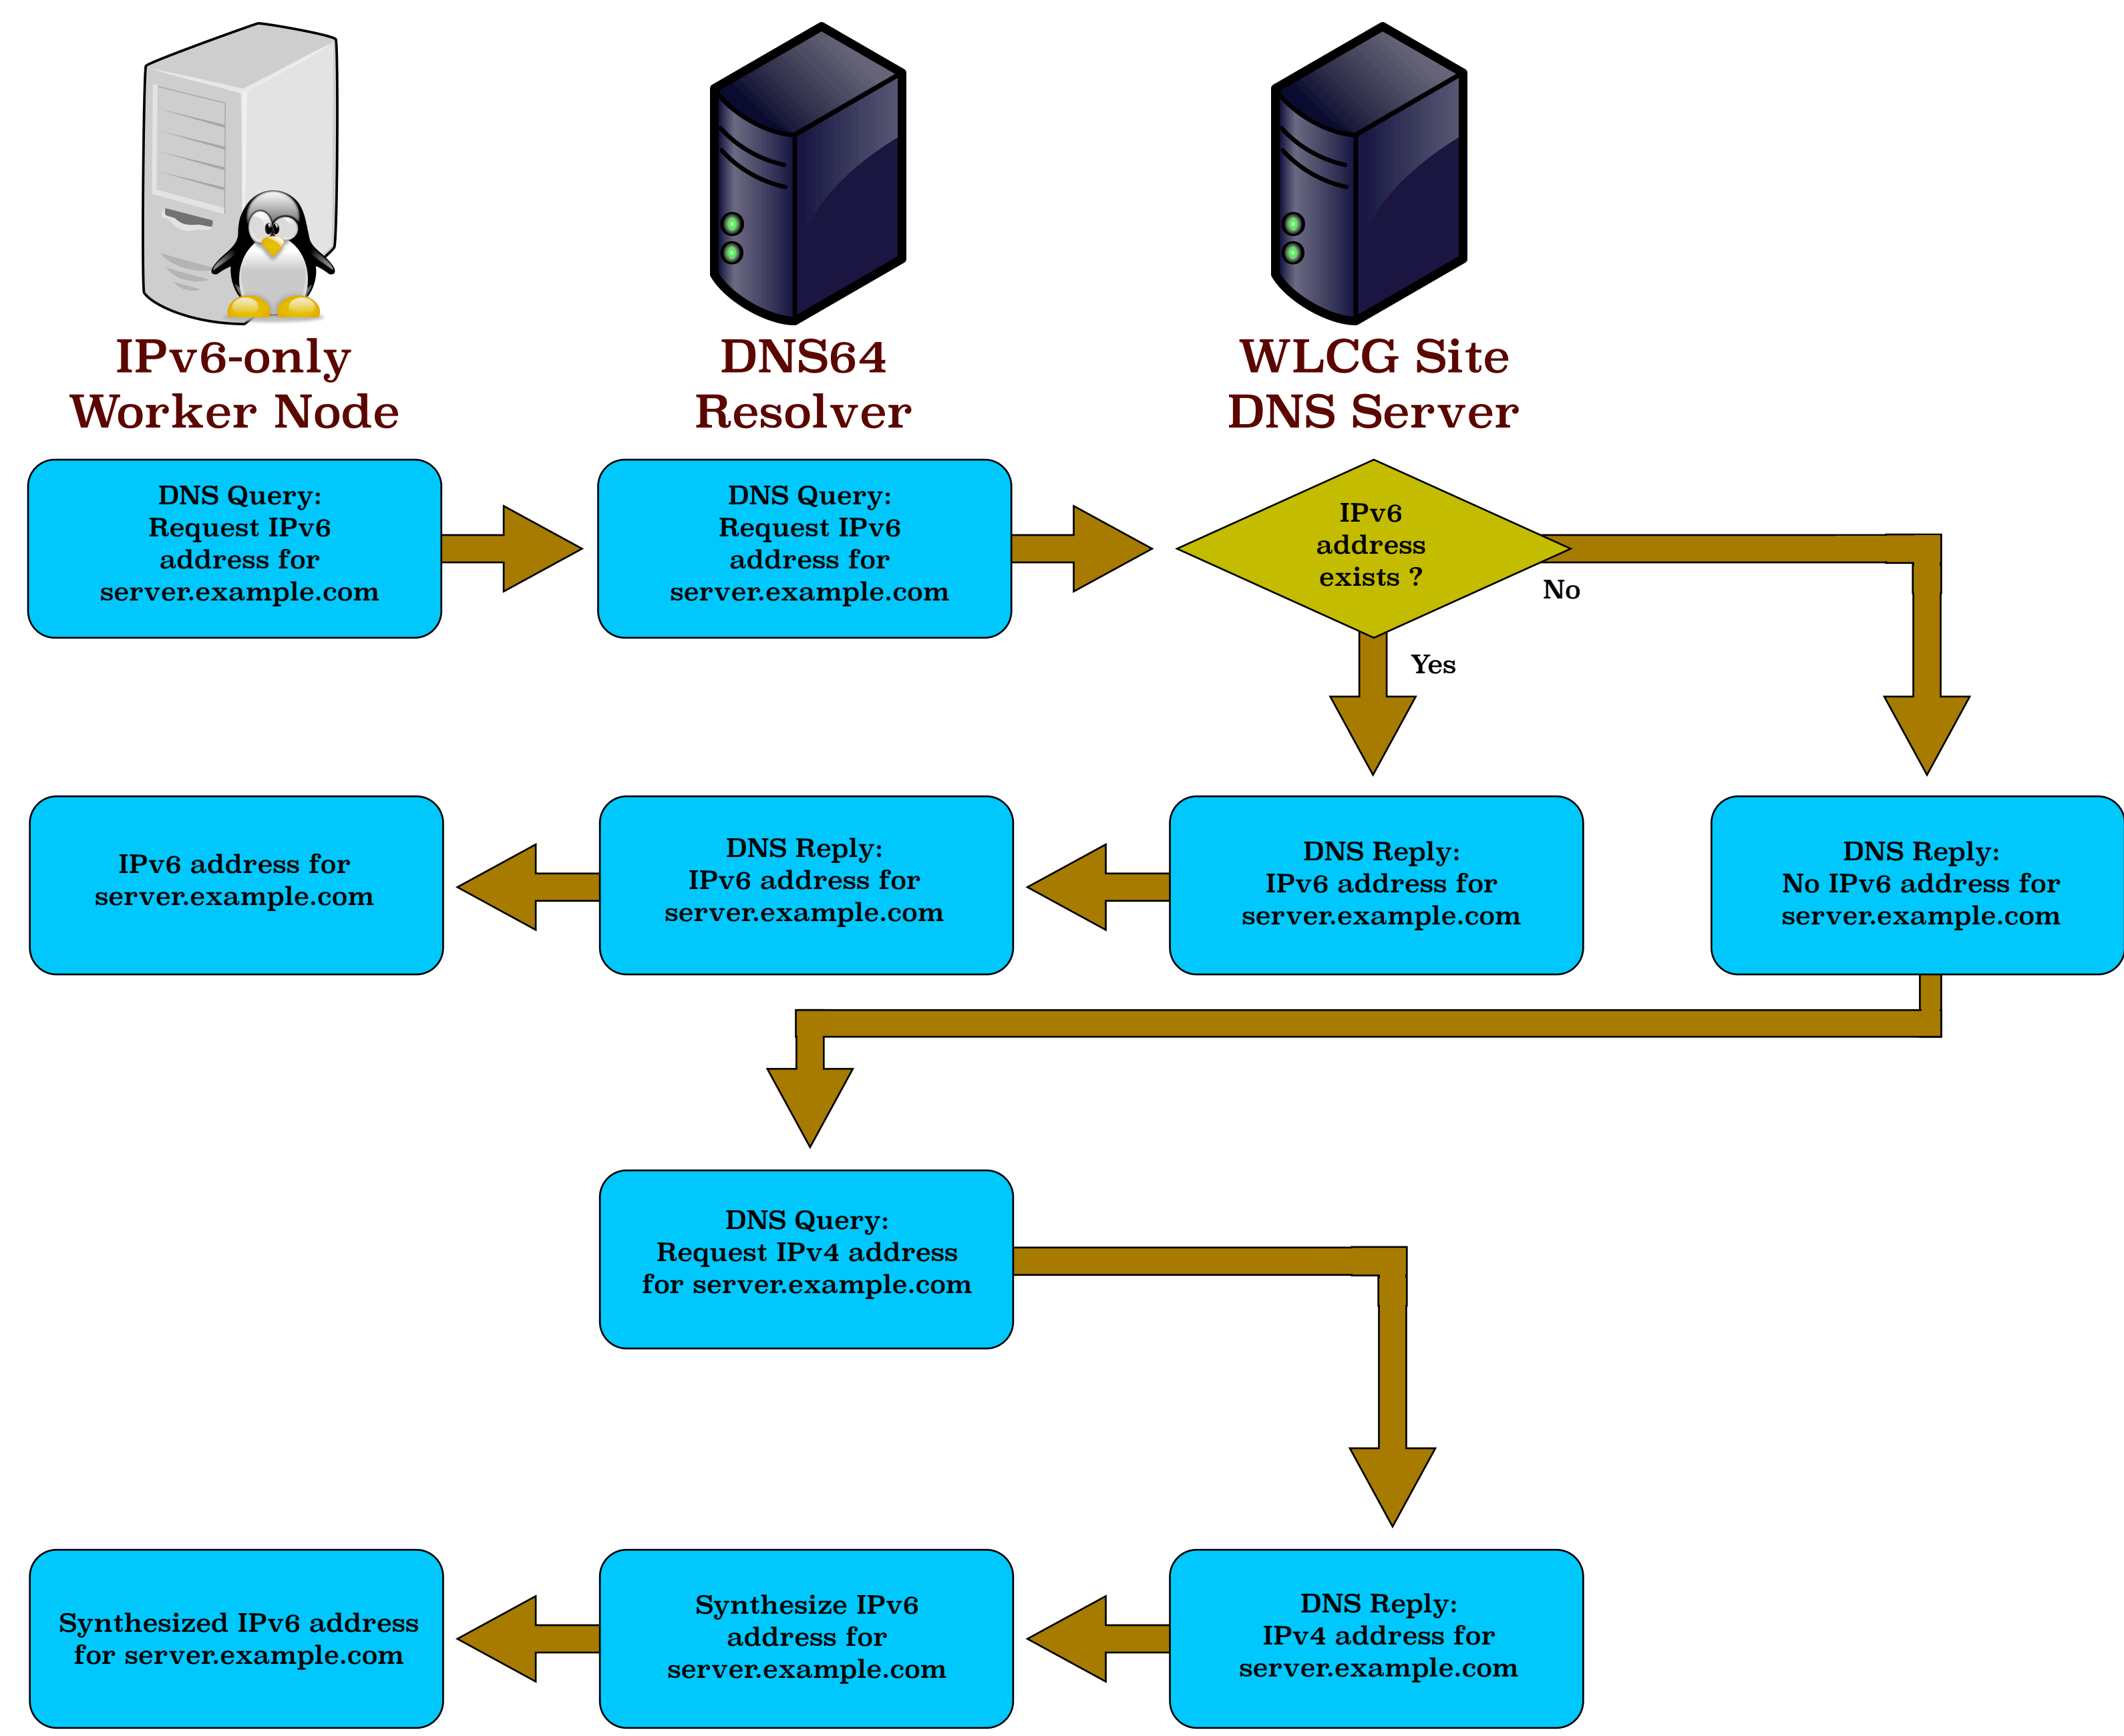
DNS64 is a mechanism which allows an IPv6-only client to always receive a valid answer to queries for AAAA records (which are the IPv6 equivalent to IPv4's A records); if the hostname has a valid AAAA record in DNS, this is passed back to the client as per standard DNS resolver behaviour.

In the absence of a valid AAAA record, the A record is queried and the 32-bit IPv4 address returned is appended to a /96 IPv6 prefix in order to form a 128-bit IPv6 address which is returned to the IPv6-only client as a synthesized AAAA record.

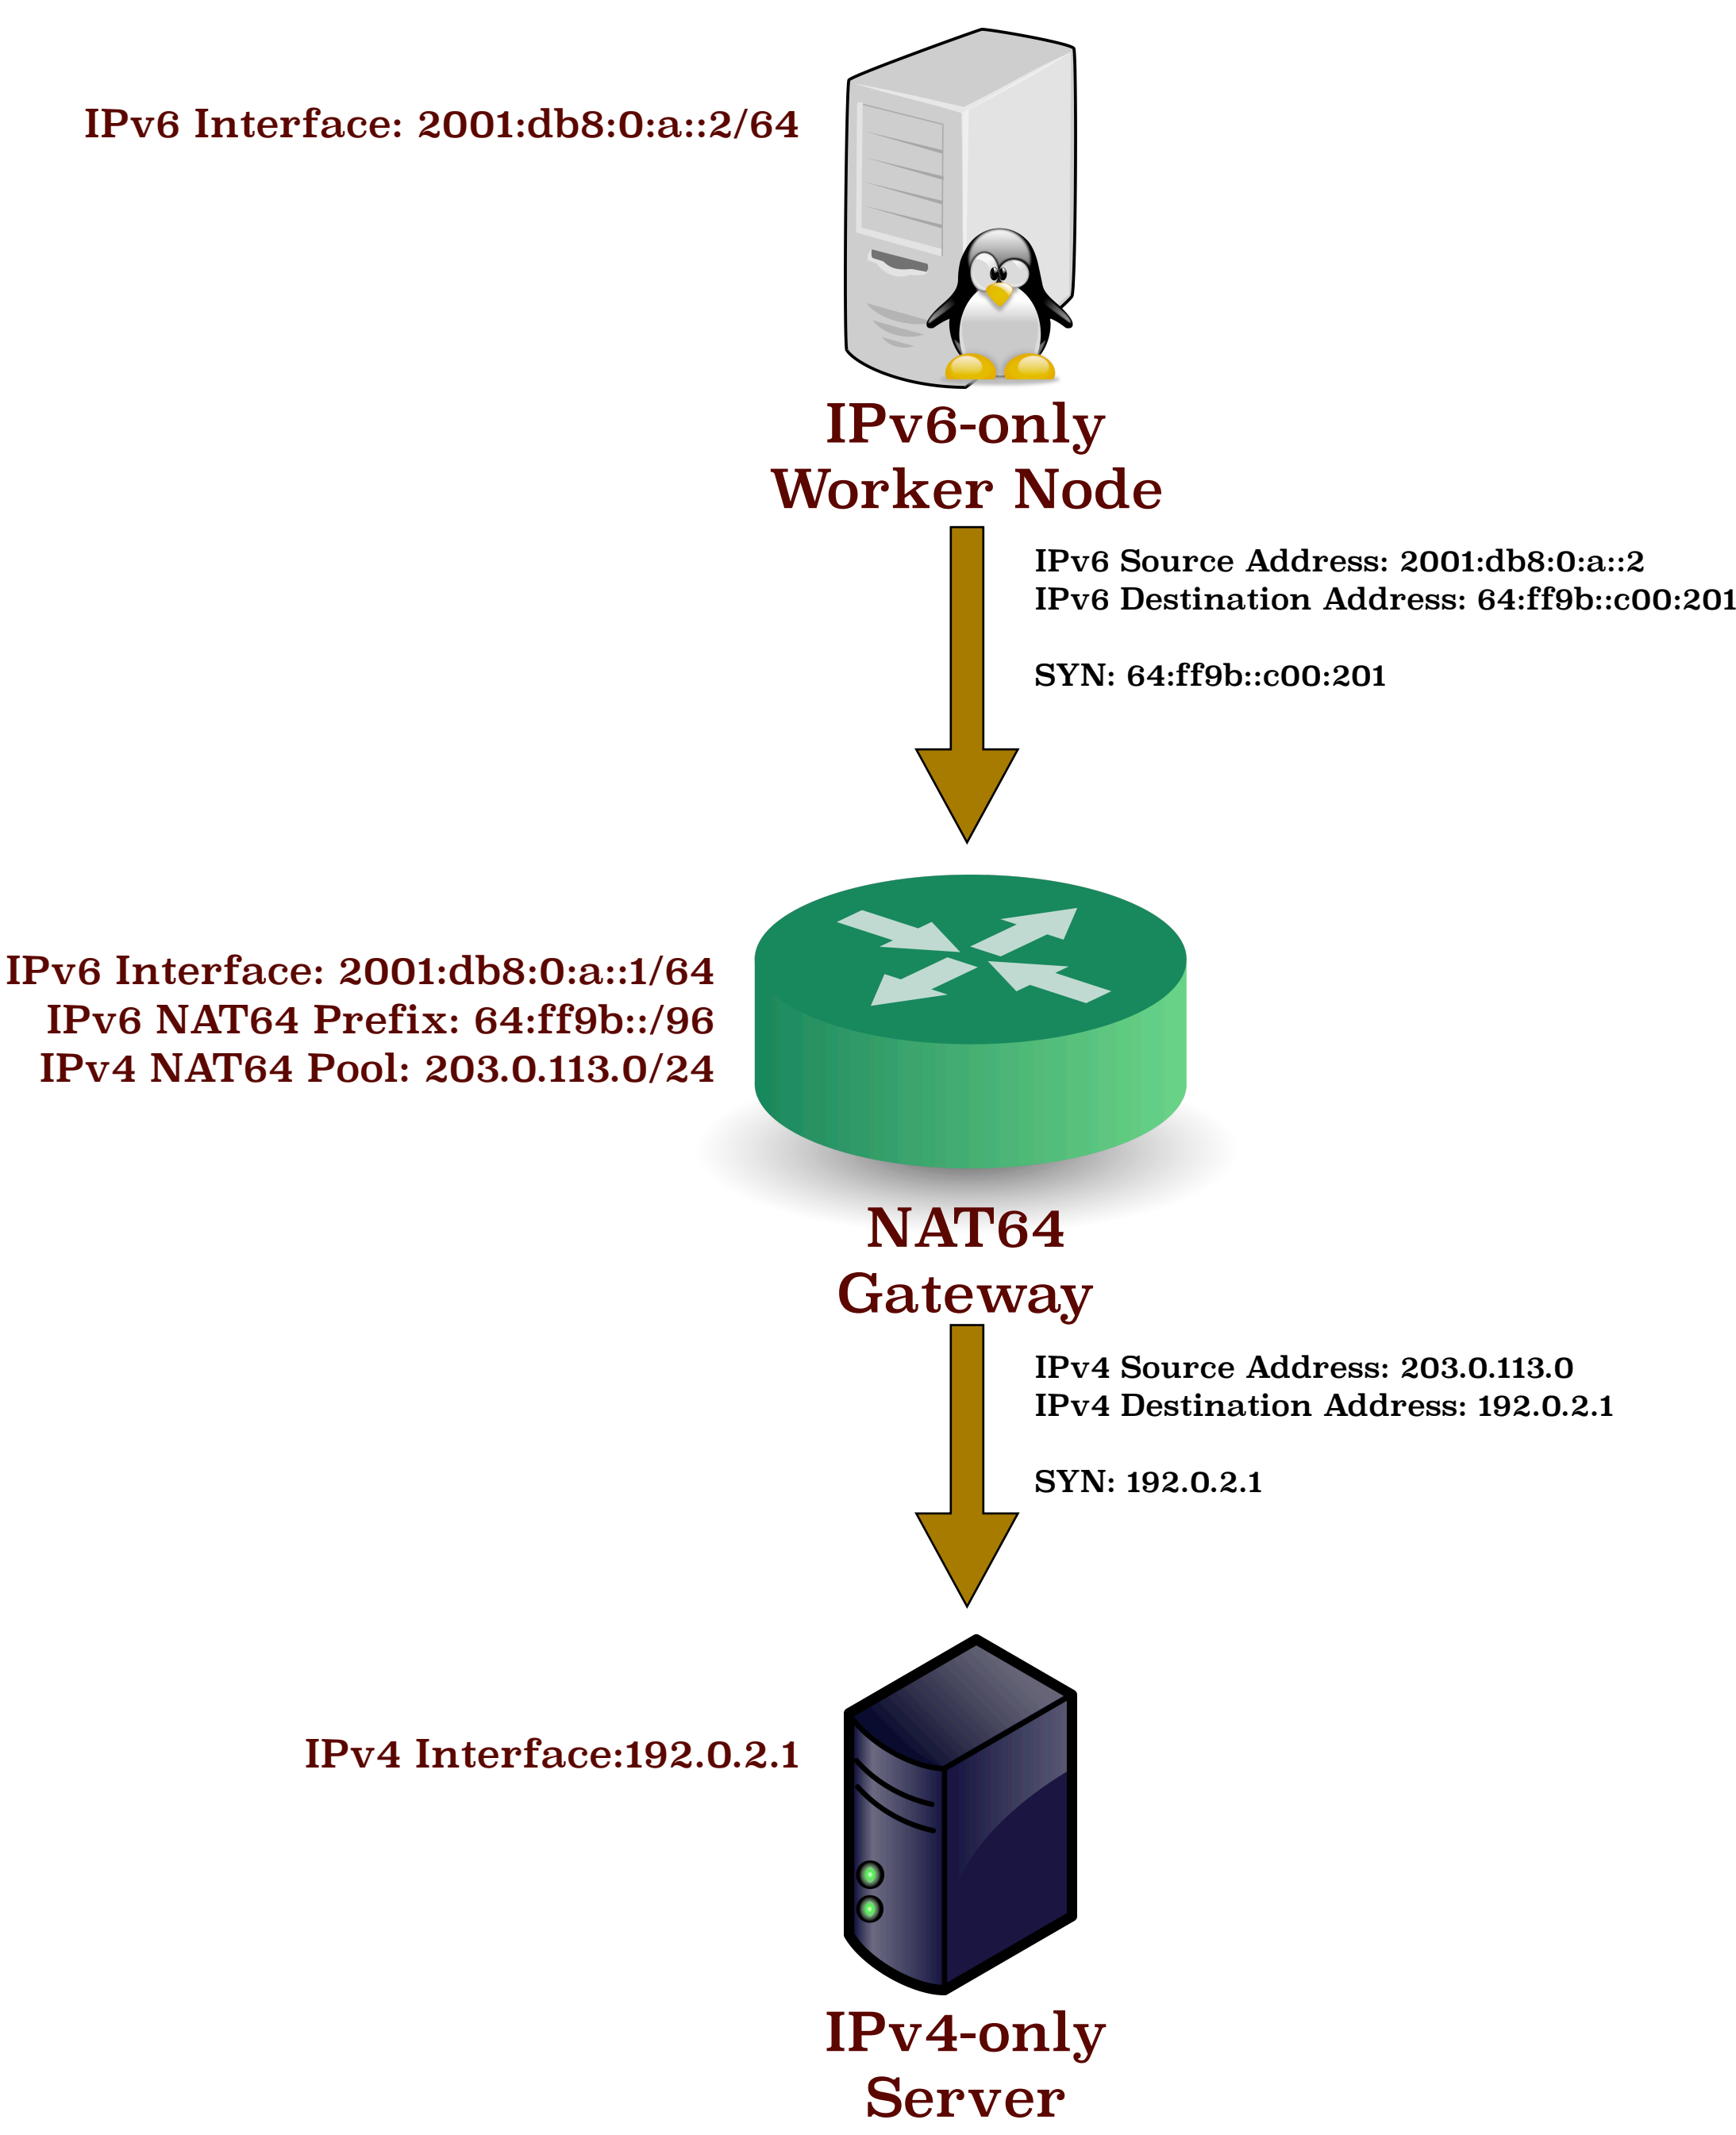
The /96 IPv6 prefix is routed towards a NAT64 gateway cluster; the client sends traffic to addresses within this prefix and the NAT64 gateway cluster extracts the last 32-bits from the destination IPv6 address and uses that as the destination IPv4 address.

The source IPv4 address used for the outbound IPv4 connection can be dynamically or statically assigned at the whim of the gateway

How DNS64 Works



How NAT64 Works



Future Plans

We are looking at the possibility of setting up anycast DNS64/NAT64 instances at a small number of WLCG sites that have high-quality and reliable IP connectivity; the use of anycast ensures that WLCG sites using these instances achieve fully-autonomous failover in the event that their nearest instance is offline (scheduled maintenance, etc)

The intention is to make it simple for smaller WLCG sites to contribute IPv6-only Worker Nodes and to entirely negate the requirement for native IPv4 connectivity; instead, shifting the administration and technical management of this service to those willing and able to do the heavy lifting.

Anycast is most commonly used to distribute DNS traffic destined for the Internet root nameservers to locally-hosted instances and serves to ensure traffic is efficiently handled; as such, anycast is considered mature and production-ready so using it to provide the same level of resilience and efficiency for WLCG sites is the next logical step in the evolution of this service.

Contacts

Terry Froy: t.froy@qmul.ac.uk

Daniel Traynor: d.traynor@qmul.ac.uk

Christopher J. Walker: c.j.walker@qmul.ac.uk

School of Physics and Astronomy, Queen Mary University of London,
Mile End Road, London, E1 4NS

Links

[DNS64 Resolver] PowerDNS Recursor: <https://www.powerdns.com/recursor.html>

[NAT64 Gateway] Jool SIIT & NAT64: <http://www.jool.mx>