THREAD GROUP

July 13, 2015

This Thread Technical white paper is provided for reference purposes only.

The full technical specification is available to Thread Group Members. To join and gain access, please follow this link: http://threadgroup.org/Join.aspx.

If you are already a member, the full specification is available in the Thread Group Portal: http://portal.threadgroup.org.

If there are questions or comments on these technical papers, please send them to help@threadgroup.org.

THREAD GROUP

Thread Border Routers

July 2015

Revision History

Revision	Date	Comments
1.0	February 10, 2015	Initial Release
2.0	July 13, 2015	Public Release

Contents

Introduction

Overview of the Border Router

In the context of a Thread Network, a **Border Router** is a device that provides connectivity of nodes in the Thread Network to other devices in external networks such as the wider Internet, local home and building IP networks, or virtual private networks (Figure 1).

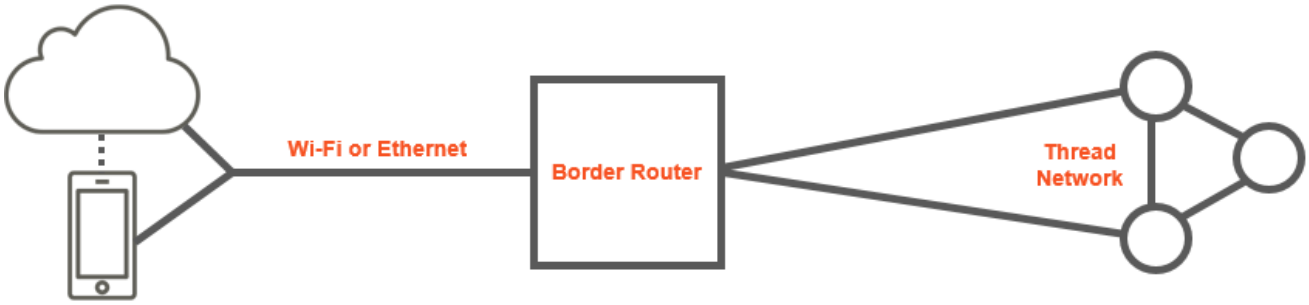

Figure 1. High Level Overview of Border Router Role

Common Characteristics

The Border Router role will be implemented by devices or products that share certain system characteristics from a networking perspective as follows.

At the **physical and link layers**, a Border Router forms a single system that includes both an IEEE 802.15.4 link-layer interface to be used for the Thread Network as well as at least a supplemental IP link-layer interface used by an exterior network (Wi-Fi or Ethernet being the most common).

At the **network layer**, a Border Router performs standard IP packet routing based on source and destination addresses contained within an IP header:

- Outgoing packets from the Thread Network interface will be forwarded to the exterior interface(s).
- Packets from exterior interface(s) will be forwarded to the Thread interface and then routed further in the Thread Network towards their end destination.
- Packet filtering or address translation may be performed based on firewall, system, or infrastructure settings.

Also at **the network layer**, a Border Router may participate in an exterior routing protocol, advertise global IPv6 prefixes and handle global scoped address allocation for nodes within the Thread Network.

At the **transport layer**, a Border Router should be transparent to end-to-end IP communication.

From a **commissioning** perspective, a Border Router will intermediate a secure, user-initiated joining of new devices to the Thread Network by means of a Commissioner device when that device is on an exterior network.

At the **application layer**, a Border Router may provide optional services, such acting as a proxy for service discovery operations on behalf of devices on the Thread Network.

Border Router Availability

Communication between devices within a Thread Network can take place without any active Border Router participating in the network. If a Border Router is not available, commissioning services must be provided by a device that participates directly in the network.

A Thread Network supports multiple active Border Routers. This has the advantage of providing redundancy and resilience, and prevents a single point of failure.

Types of Border Router Devices

There are two categories of products that may implement a Thread Border Router:

- Consumer premises networking equipment and residential gateways
- Consumer products that include a Thread interface as well as alternative connectivity

In a typical home network, the most common scenario is for specialized networking devices such as access points or home routers to provide routing and Internet access services between the respective local area network and an exterior WAN (Wide Area Network). These devices are usually referred to as being Customer Edge or Customer Premises Equipment from the perspective of the provider.

This category of specialized network equipment or on-premises gateways may also be provisioned with a Thread physical network interface that will allow these devices to fulfill a Thread Border Router role.

However, Thread Border Router functionality can also be easily included within consumer home products such as simpler appliances that include both Thread and Wi-Fi interfaces.

Multiple Link Layer Interfaces

A pre-condition of implementing a Thread Border Router is the availability of multiple link layer interfaces.

Figure 2 and Figure 3 illustrate a system overview of the two categories of Border Router devices from the perspective of on-board interfaces.

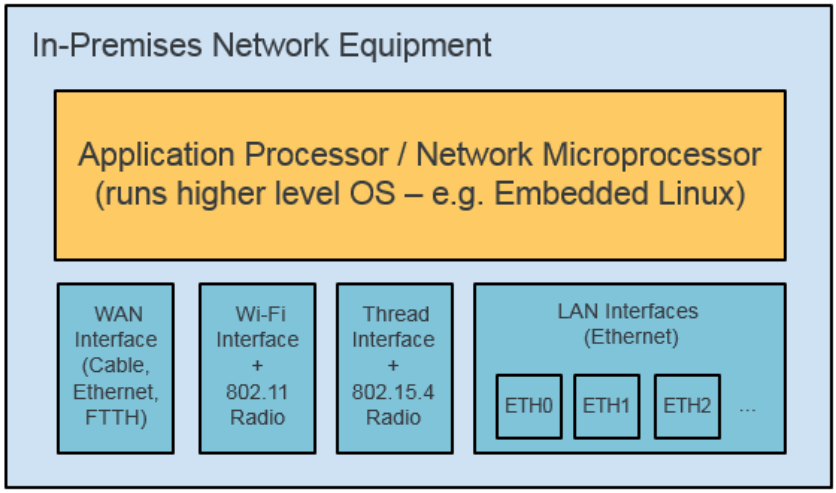

Figure 2. Border Router as In-Premises Networking Device

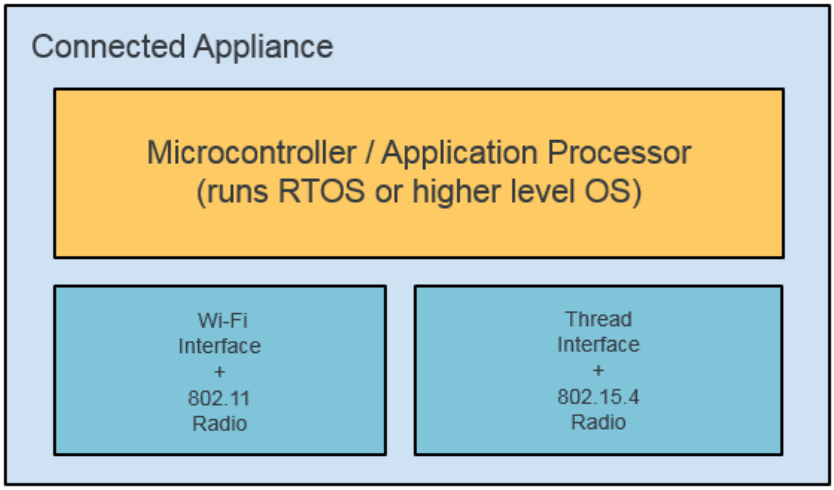

Figure 3. Border Router as Consumer Appliance

It is expected that networking equipment will include a Thread interface as a specific wireless LAN port similar to Wi-Fi and provide the possibility of IP routing between multiple internal interfaces to the WAN and/or a VPN (Virtual Private Network) connection.

Consumer home products will usually be provisioned with a lower number of interfaces and usually run within a more constrained system.

Network Layer

Network Layer Overview

The ability to forward packets to and from an external network link, and also to participate in an exterior routing method or protocol is the main characteristic of the Border Router. This white paper covers the fundamentals of network layer routing to and from an external network by the Thread Network.

IPv6 Global Addresses

A Thread Network only operates using IPv6. The use of IPv4 addressing is not supported for communication within the Thread Network.

Individual Thread Network devices support participation in the global IPv6 (Internet Protocol version 6) infrastructure, such as being part of the IPv6 Internet. This is achieved by means of GUAs (Global Unicast Addresses) that are described in [RFC 4291].

Each Thread node can be assigned at least one GUA when the upstream infrastructure for delegating a global prefix via a Border Router is available.

The Border Router notifies information on the global prefixes it serves to the Thread Leader, which adds it to a Network Dataset, and then distributes it within the Thread Network.

In some cases, the Border Router may also handle individual global address assignment to Thread nodes by means of DHCPv6 (Dynamic Host Configuration Protocol version 6) messages described in [RFC 3315]. The option for nodes to use SLAAC (Stateless Address Autoconfiguration) addresses based on a prefix advertised by the Border Router is also available.

Thread Border Routers can obtain global prefix assignments by participating in an exterior prefix distribution protocol such as DHCPv6-PD, L2TP-VPN, or HNCP (HomeNet Control Protocol). If necessary, they may participate in exterior routing domains with a routing protocol such as RIP

(Routing Information Protocol), OSPF (Open Shortest Path First), IS-IS (Intermediate System-to-Intermediate System, and others.

IPv6 Unique Local Addresses

Thread Network devices support ULAs (Unique Local IPv6 Unicast Addresses) that are described in [RFC 4193].

A Thread Network uses a specific category of ULAs for the purposes of mesh routing and management within the network. In the context of the Thread Network, these are called MLAs (Mesh Local Addresses)–either ML-EID (Mesh-Local Endpoint Identifier) or ML-RLOC (Mesh-Local Routing Locator)–and are identified by a ULA prefix referred to as the MLP (Mesh Local Prefix).

Communication using MLAs is not meant to be routable to the exterior of a Thread Network via a Border Router.

Supplemental ULA prefixes MAY be used to assign other ULAs to Thread interfaces. This allows creation of IPv6 fabrics spanning multiple on-premises site-local subnets and wide-area virtual private networks.

Such ULAs are useful when either upstream WAN (Wide Area Network) infrastructure does not provide means for IPv6 global prefix delegation or when the application use case either does not need or specifically precludes routing on the Internet.

Assignment of Supplementary ULA prefixes is handled identically to global prefixes from the perspective of how they are provisioned to the Thread Network by the Border Router.

Supplementary site-local ULA prefixes may be generated by Customer Edge routers adhering to recommendations in [RFC 7084].

Assignment of Supplementary ULA prefixes is handled identically to the global prefixes from the perspective of the Thread interface on the Border Router and that of the global scoped information in the Network Data distributed by the Thread Leader.

Notification and Propagation of Network Data

Thread interior networks do not use the IPv6 Neighbor Discovery protocol. Global prefixes and Supplementary ULA prefixes are not distributed using Router Advertisement messages, and addresses are not assigned using stateless auto-configuration. Instead, prefixes are advertised in Network Data messages from the Thread Leader.

The TMF (Thread Management Framework) protocol is used for notification of Network Data from the Border Routers and DHCP servers to the Leader. TMF is based on CoAP (Constrained Application Protocol) that is described in [RFC 7252].

When forwarding packets with an exterior destination originating from an interior node, regular IP mesh routing within the Thread Network is used to reach the Border Router providing an external route for the prefix as advertised in the Network Data.

Some Thread Child end devices may receive a unicast containing all or a subset of the Network Data as transmitted by their Parent. A constrained Child end device, such as a sleeping node, can choose to maintain only a subset of the Thread Network data that is considered more stable and not receive temporary data which has a limited lifetime.

Both stable and temporary instances of the Thread Network data are versioned by the Leader with the version being used by Parent to update end nodes.

The Thread Network data includes:

- **On-Mesh Prefix Set** contains the IPv6 prefixes available to nodes on the Thread Network, along with the Border Routers which provide them and preference indication on address assignment using DHCPv6 or SLAAC
- **External Route Set** contains information on the available external routes for packets originating in the Thread Network with an exterior destination.
- **6LoWPAN Context ID Set** consists in the 6LoWPAN context information used to compress the size of the global addresses. Detailed information of how the Context ID is used for 6LoWPAN encoding can be found in the **"Thread Usage of 6LoWPAN"** white paper.
- **Server Set** contains information on the available servers providing standard or vendor-specific network services to Thread nodes.

When forwarding packets with an exterior destination from an interior node, regular IP mesh routing within the Thread Network detailed in the **"Thread Stack Fundamentals"** white paper is used to reach the closest Border Router providing an external route for the prefix as advertised in the Thread Network data.

Figure 4 illustrates the notification and propagation of Thread Network Data and GUA assignments using DHCPv6.

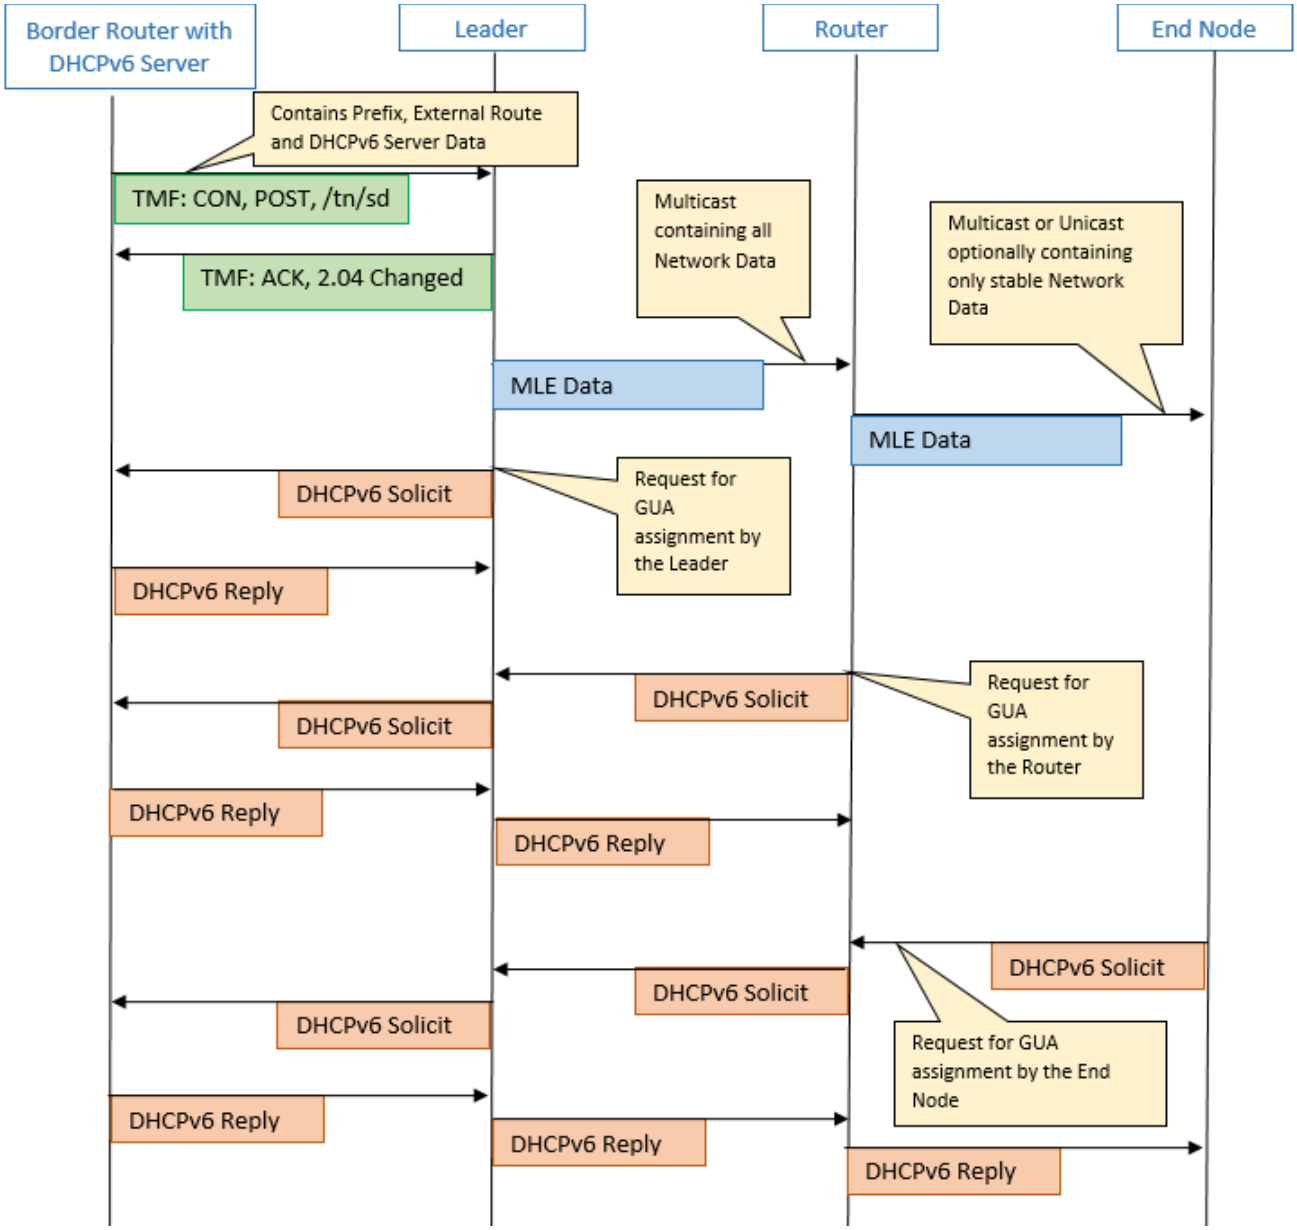

Figure 4. Notification and Propagation of Thread Network Data and GUA Assignments using DHCPv6

Coping with IPv4 and Hybrid Infrastructure

As of 2015, IPv6 deployment in Internet infrastructure has seen advances. However, usage within residential networks is still perceived to have significant limitations, especially regarding capabilities and provisioning of on-premises equipment for sub-netting and prefix distribution.

While some ISP capabilities for native IPv6 infrastructure are expanding rapidly, there is a large existing base of Customer Edge and Customer Premises devices, such as home and small business routers, modems, or access points for which IPv6 functionality is either not available or otherwise complex to configure.

It is expected that from the perspective of communication between Thread devices and exterior networks, coexistence and coping with IPv4-only local infrastructure or with limitations in IPv6 deployment in residential networks will be accomplished primarily by capabilities of the Border Routers.

The most common cases for home network deployments is where a device or home appliance which can act as a Thread Border Router can be assigned only a site-local IPv4 address, or in certain cases only a global /128 IPv6 address with limited capabilities for downstream prefix delegation.

While less desirable than usage of natively provisioned GUAs and ULAs, the Thread Border Routers are expected to revert in this case to implementing a form of NAT (Network Address Translation) so that the application use case is still achieved.

In particular, for communication between the interior IPv6 Thread Network and either a residential IPv4 site-local network or the IPv4 Internet, Thread Border Routers may use a variant of NAT64, within the frameworks specified by [RFC 6144] and [RFC 6052].

However, in all situations, communication and external prefix packet routing within the Thread Network still operates as described in the previous sections.

Transport and Application Layer

Packet Filtering and Port Forwarding

Border Routers are expected to provide ingress and egress traffic filtering capabilities and offer protection to the Thread interior network based on recommendations in [RFC 4874] and [RFC 6092].

Given the limited bandwidth and power constraints of most categories of Thread nodes, the Border Router role in managing packet filtering is essential to avoid intentional or accidental

flooding or DoS (Denial of Service) within the Thread mesh based on traffic originating from the exterior.

Beyond firewalling for security attacks, Thread Border Routers are expected to provide rate limiting or ultimately deny forwarding of otherwise legitimate ingress application traffic when it is detected that excessive unicast or multicast data flow may be disruptive for the constrained interior network.

However, such rules should not preclude end-to-end communication of Thread devices with devices on the exterior networks, especially for outbound traffic.

If a Border Router is itself a constrained or battery-powered device, it will attempt to delegate port filtering and forwarding rule configuration to an upstream device.

Seamless configuration of port control and filtering between Thread nodes and the exterior can be optionally achieved using a lightweight control protocol such PCP described in [RFC 6887]. Beyond port control capabilities, these protocols can also be used to configure NAT services when such translation is needed for external network communication.

Role in Commissioning

Border Routers play an essential role in Commissioning of new devices to the Thread Network. In this regard, they act as a relay for messages between external Commissioners (such as a mobile device) and devices on the Thread Network, such as the Leader which provides arbitration for petitioning of multiple potential Commissioners, or a factory new device for which joining was initiated.

The role played by the Border Router in Commissioning is described in further detail in the **"Thread Commissioning"** white paper.

Example of Border Router Operation and Role Fulfillment

Functional Roles of a Border Router

In the context of operation within the interior Thread Network, the Border Router can provide a set of discrete services and functions. Some are related to the multiple interface capability, with others being orthogonal to it.

Some services are optional at runtime because the network conditions or system configuration can cause to dynamically enable or disable these functions and services. Table 1 enumerates them, along with events which can determine activation or de-activation of functions.

Table 1. Functional Roles of the Border Router

Border Router Functional Role	Role Access Conditions
External Routing (including default route)	Role can be enabled, disabled, or configured by a high-level application on the Border Router. Disabled when external interface loses connectivity.
Thread Interior Routing	Managed autonomously by the Thread Network stack.
Server for external prefix Network Data	Managed by a high-level application on the Border Router and operated by the Thread Network stack. Multiple Border Routers can synchronize on this role using the exterior network.
Thread Leader	Managed autonomously by the Thread Network stack (orthogonal to External Routing role).
Commissioner	Managed by a high-level application on the Border Router. Must petition Leader and get favorable outcome to enter role (orthogonal to External Routing role).
Joiner Router	Managed autonomously by the Thread Network stack.
Commissioning Relay	Managed by the Thread Network stack. Configured by a high-level application.
Service Discovery Server	Managed by a high-level application on the Border Router.

Border Router Role Transitions

The following sections detail a typical Thread Network formation and subsequent operation from the perspective of a Border Router. Initially, several functions and roles are collapsed to the same device which is used to form a new Thread Network and then acts as a Leader. As more devices join the network and the user chooses exterior Commissioners, some of the roles become more distributed.

Network Formation by the Border Router as Leader

Figure 5. BR1 starts a Thread Network as Leader, activates external interface

Node BR1 has both Thread and an external network interfaces available and can act as a Border Router. However, as it is essentially the single device in the Thread Network, minimal border functions are active at this point.

It is expected that the exterior interface becomes configured and provisioned for access to the home network (LAN or WLAN) and the Internet. For simplicity, this scenario assumes the exterior interface is also used to auto-configure delegation of a globally scoped /64 IPv6 prefix for node address assignment on the Thread Network.

Border Router exercises roles: Leader, Server for global prefix

Joining with Co-located Commissioner

Figure 6. BR1 acts as an internal Commissioner for joining R1

A user initiates joining of a new node (R1) to the network. R1 has only a Thread interface available for communication. Assuming BR1 is provided with an adequate UI, it may also be used to act as a Commissioner. The commissioning session is native to the Thread Network interfaces and no data flow across the Thread Network border is necessary to finalize commissioning and joining of R1.

Border Router exercises roles: Leader, Joiner Router, and Commissioner

Address Assignment, External Routing

After R1 joins the Thread Network, it becomes a Router and will start receiving multicast routing and Network Data MLE advertisements from BR1 acting as a Leader. The Leader and a Global prefix DHCPv6 server roles are collapsed on BR1. The information in the incoming

advertisements is subsequently used by R1 as a DHCPv6 client to unicast solicit messages to BR1 for assignment of Global addresses.

R1 can subsequently use the globally scoped addressing and the external route via BR1 to interact with nodes on the home WLAN or the Internet at the application layer.

Border Router exercises roles: Leader, External Routing, DHCPv6 Server for global prefix, Packet and Port Filtering

Petitioning by External Commissioner

Figure 7. External Commissioner registers with BR1

A user initiates petitioning for a new Commissioner (C) not participating in the Thread Network. The new Commissioner can be a mobile device connected to the same WLAN network segment or which can otherwise establish a connection to the exterior interface of BR1.

A shared petitioning passphrase is used to complete a DTLS handshake between C and BR1 which opens a Commissioning session between the nodes. Because BR1 also exercises the Leader role, the message flow for petitioning takes place only on the external network.

This scenario assumes BR1 successfully registers the new device as Commissioner in its Leader role. BR1 will subsequently advertise itself within the Thread Network as a commissioning relay.

It is possible for the alternate Commissioner role to be assumed via petitioning by another device participating in the Thread Network that is different from the BR1, such as the newly joined router R1. For an analysis of this use case, see the "**Thread Commissioning"** white paper.

Border Router exercises roles: Leader, External Routing, External Petitioning Server, Commissioning Relay

Border Router delegates roles: Commissioner

Joining with External Commissioner

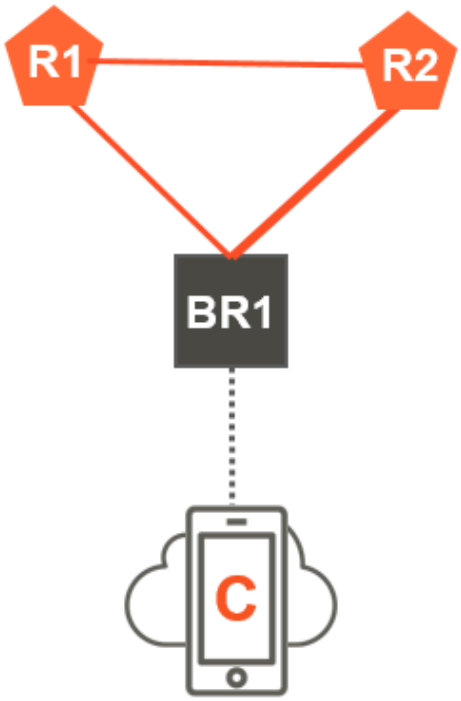

Figure 8. BR1 relays commissioning for C to allow R2 to join

A user initiates joining of a new node (R2) to the network. R2 has only a Thread interface available for communication. External Commissioner (C) is used after C completes petitioning with BR1 as described above.

Based on Thread Network topology, either BR1 or R1 can be selected as a Joiner Router. In the simplest case when BR1 is chosen, an initial DTLS handshake is sent directly from R2 to BR1. Otherwise, if R1 is chosen as the Joiner Router, the DTLS handshake between R2 and BR1 is routed across R1 by means of TMF messages within the secured Thread Network.

BR1 further relays the DTLS handshake from the Joiner to the Commissioner C by means of messages exchanged across the secure commissioning session established after petitioning on the external network.

If the commissioner to joiner handshake completes successfully and BR1 is the Joiner Router, BR1 will use the key resulted after commissioning to securely distribute network parameters and security material to Joiner R2.

Once R2 joins, it will also receive Network Data advertisements from BR1 acting as a Leader and use DHCPv6 to get global scoped address assignment.

Border Router exercises roles: External Routing, Leader, Commissioning Relay, Joiner Router

Border Router delegates roles: Commissioner

Transitioning from the Leader Role

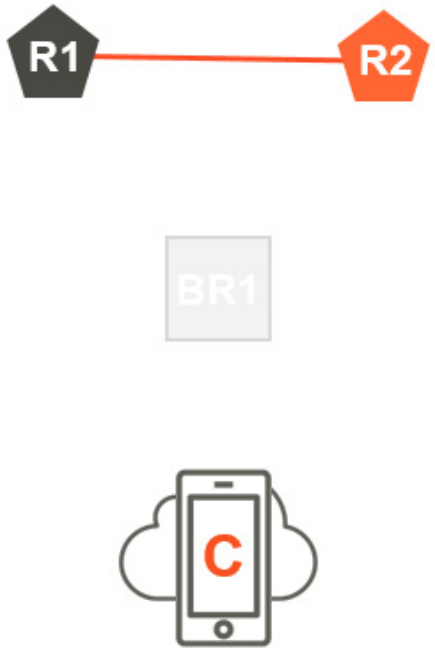

Figure 9. BR1 becomes unavailable. R1 takes over Leader role.

A user may temporarily power off device BR1 (for example, for maintenance purposes). As BR1 is also the Leader, when it is no longer be active in the Thread Network for a defined time interval, another router will take over the Leader role.

We can assume R1 becomes the new Leader. When a new Leader starts, the Thread Network Data is reset. Without a Border Router active in the network, the new Thread Network Data advertised by R1 also no longer contains information for external routing.

Powering BR1 back on results in BR1 acknowledging R1 as the new Leader and subsequently BR1 uses TMF to dispatch its external route and DHCPv6 information Thread Network Data. R1 starts propagating the Thread Network Data advertising BR1 to the network. Routing to the exterior is restored.

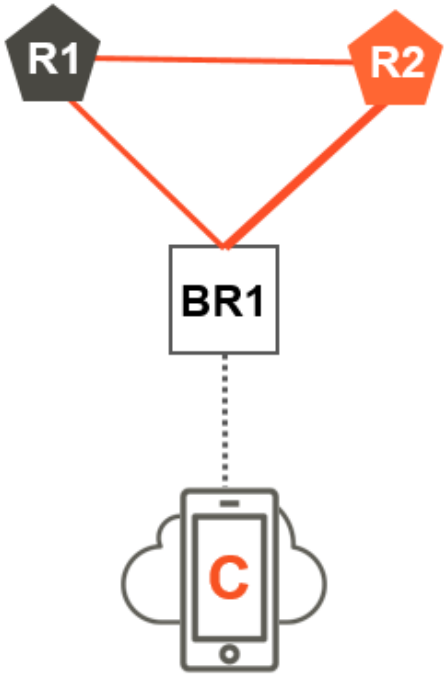

Figure 10. BR1 returns to the network. R1 remains Leader.

If joining a new device to the Thread Network becomes necessary, a new petitioning session needs to establish a Commissioner (C). The new Leader R1 is used for the petitioning arbitration between multiple potential Commissioners.

TMF messages are used by BR1 to interact with the Leader R1. BR1 can acts either on behalf of Commissioners on the external network, or to register itself as a Commissioner. Commissioning then proceeds to occur similarly to as described in the previous sections.

Border Router exercises roles: External Routing, Commissioning Relay, DHCPv6 Server for global prefix

Border Router exercises roles: Leader, Commissioner

Delegating the DHCPv6 Server Role

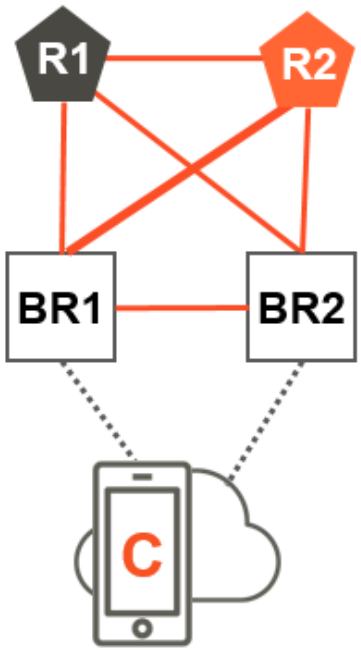

Figure 11. BR2 joins the network and takes over DHCPv6 role. BR1 and BR2 provide external routing.

Assume a new device BR2 with border routing capabilities joins the network. BR2 is connected to the same external LAN segment as BR1, or otherwise BR2 can reach BR1 on the exterior network. One can also assume BR1 and BR2 have been delegated the same global prefix by their upstream configuration. It is expected that BR1 and BR2 will use an election protocol that takes place on the external network to establish which of the Border Routers will run an authoritative DHCPv6 server instance for address assignment based on the global prefix within the Thread Network.

Both BR1 and BR2 will use TMF to send their information to the Leader for Thread Network Data to include them as part of the External Route set for the global prefix; however, only one of the devices will provide DHCPv6 server information.

BR1 exercises roles: External Routing, Commissioning Relay

BR1 delegates roles: Leader, Commissioner, DHCPv6 Server for global prefix

Other Topics

IPv6 Transition Technologies

In the context of transitioning from IPv4 to IPv6, several temporary mechanisms that facilitate this transition have become relatively widespread.

As mentioned in **Coping with IPv4 and Hybrid Infrastructure**, a variant of NAT64 may be needed for communication from a Thread Network via the Border Router to an IPv4-only on-premise LAN or the IPv4 Internet and is recommended to be implemented on Border Routers.

While Border Router applications can choose to integrate other transition technologies when IPv6 infrastructure is not available in the upstream, most non-native mechanisms have a setup complexity that is not expected to be dealt with directly by residential users. Border router applications are expected to provide autonomous configuration of such mechanisms or refrain from using them.

Homenet

The Homenet workgroup within the IETF was formed to resolve the issue of current state of high configuration overhead as residential networks become more complex.

From the perspective of integrating Thread Networks in a Homenet environment, it is expected that specification work done in Homenet can be re-used for more autonomous configurations of network interfaces for Thread Border Routers and may allow Thread Border Routers to synchronize their operation more seamlessly on the external network. In particular, HNCP (Home Networking Control Protocol) is meant to enable automated configuration of addresses, network borders and the seamless use of routing protocols within a residential network fabric.

References

Document	Title
[RFC 2661]	Layer Two Tunneling Protocol "L2TP"
[RFC 3315]	Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
[RFC 3633]	IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6
[RFC 4291]	IP Version 6 Addressing Architecture

Document	Title
[RFC 4864]	Local Network Protection for IPv6
[RFC 4941]	Privacy Extensions for Stateless Address Autoconfiguration in IPv6
[RFC 6144]	Framework for IPv4/IPv6 Translation
[RFC 6052]	IPv6 Addressing of IPv4/IPv6 Translators
[RFC 6092]	Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service
[RFC 6887]	Port Control Protocol (PCP)
[RFC 7084]	Basic Requirements for IPv6 Customer Edge Routers
[RFC 7252]	https://tools.ietf.org/html/rfc7252
[RFC 7368]	IPv6 Home Networking Architecture Principles
[draft-ietf-homenet-hncp]	Home Networking Control Protocol