Chapter 10
Routing and Addressing

Yingzhen Qu, Adrian Perrig, and Daniel King

10.1 Introduction

The current Internet, which has evolved for more than 50 years, is facing a set of unique challenges, both technically and commercially. The exponential growth of the Internet and emerging demands from connected devices, increased mobility, security and resilience are met through incremental updates. Routing protocols have been critical networking technologies, and continuous development of routing protocols is essential to provide network services, which are the building blocks for new applications and services.

Figure 10.1 classifies widely used routing protocols into different categories.

Distance vector protocols are based on the Bellman-Ford algorithm, and are also referred to as routing by rumor, as they rely on neighbor-based information. Routers iteratively calculate the best routes to others as routing information propagates through the network. Common distance vector protocols include Enhanced Interior Gateway Routing Protocol (EIGRP) (https://tools.ietf.org/html/rfc7868) and Routing Information Protocol (RIP) (https://tools.ietf.org/html/rfc2453).

In link state protocols, each router floods its connectivity information to all other routers and locally calculates the shortest paths to them using Dijkstra's algorithm. Any change of link status (e.g., an interface shutdown) will be advertised to all

Y. Qu (✉)
Futurewei Technologies Inc., Santa Clara, CA, USA
e-mail: yingzhen.qu@futurewei.com

A. Perrig
Department of Computer Science, Network Security Group, ETH Zurich, Zürich, Switzerland
e-mail: adrian.perrig@inf.ethz.ch

D. King
Department of Computing and Communications, Lancaster University, Lancaster, UK
e-mail: d.king@lancaster.ac.uk

M. Toy (ed.), *Future Networks, Services and Management*,
https://doi.org/10.1007/978-3-030-81961-3_10

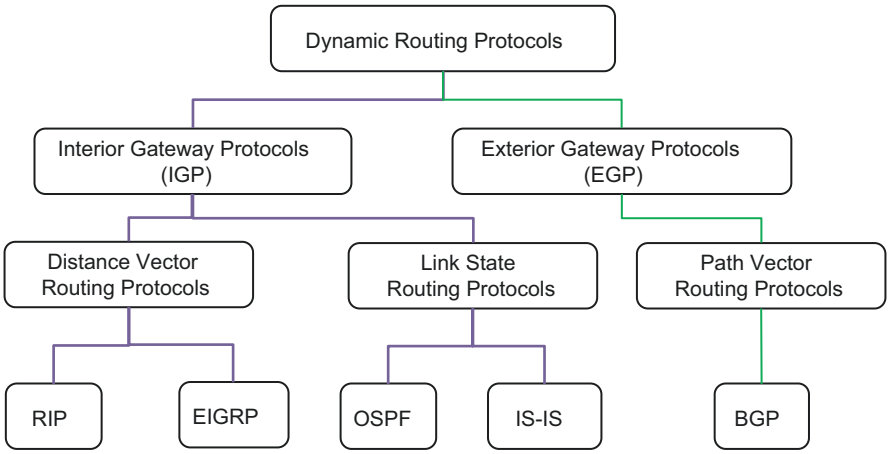

Fig. 10.1 Routing protocol category

routers in the network, allowing them to recalculate the shortest paths and maintain an up-to-date view of the entire network topology. Examples of link-state protocols include Open Shortest Path First (OSPF) (https://tools.ietf.org/html/rfc2328) and Intermediate System to Intermediate System (IS-IS) [1].

A path vector protocol, such as the Border Gateway Protocol (BGP) [2], iteratively builds up an Autonomous System (AS) path for each destination network. An autonomous system is composed of a set of routers under a single entity's administrative control. One of the advantages of path vector protocols is each destination network has a path dynamically added to it. Therefore, a loop is detected if the AS finds its own AS number in the path received. Routing protocols can be broken down just one more level as interior gateway protocol vs. exterior gateway protocol. Interior gateway protocols are used within an autonomous system, typically not routed between autonomous systems from the ground up, and interior gateway protocols are designed to fast route convergence, e.g., how fast to route around a link failure in a network. Exterior gateway protocols are used to route traffic between autonomous systems from the ground up, and they were designed to hold large amounts of routes, e.g., the routing table of the Internet. Another important thing is the ability to perform routing policies. For example, if there are two Internet providers, a routing policy can be defined for given prefix to prefer an ISP over another. This is done by allowing a preferred ISP ingress into the autonomous system. To access external resources to the autonomous system, the autonomous system needs to build a neighborship with another autonomous system. This is how routing on the Internet works. Of course, there are way more complicated things that make the Internet work. But at the networking level, fundamentally, that's how the Internet is shared and how it works.

Recent research and investigation for the future of the Internet identified several technology objectives, including contextual addressing, application-aware

networking, increased stability and security, faster convergence, and decreased operational costs.

At the core of the Internet are routing protocols, including OSPF, IS-IS, and BGP, facilitating how Internet routers communicate with each other to distribute information that enables them to select routes for Internet connectivity. Existing routing protocols likely need to be enhanced. New routing protocols may also be required to meet the new requirements for the emerging requirements and long-term Internet evolution goals.

10.2 Addressing

Internet Protocol (IP) addressing facilitates how one device attached to the Internet is distinguished from every other device. They are used to direct requests to an appropriate destination (destination address) and indicate where replies should be sent (source address). Due to the rapid growth of the Internet and exponential increase of connected devices, several short-term fixes have been developed for coping with Internet addressing demands. The continued growth and deployment of the Internet of Things (IoTs) and new network types such as space-networking will place new requirements on existing addressing schemes.

10.2.1 IP Address

An Internet Protocol address (IP address) is a number assigned to each device connected to a network using the Internet Protocol (IP). An IP address is used to both identify a host and the location of the host.

There are two versions of the Internet Protocol commonly used today. The original version is Internet Protocol version 4 (IPv4), which was deployed on the ARPANET [3] in 1983 and still carries the most traffic on the Internet today. An IPv4 address is a 32-bit number and is typically written in dot-decimal notation, such as 192.168.1.1. As the number of devices on the Internet increases, the IPv4 addresses have been exhausted at the IANA level since 2011. A new version of IP (IPv6) using 128-bit number was standardized in 1998 [4], and the deployment of IPv6 started in the mid-2000s.

10.2.2 Name-Based Network

Compared with using IP address as both the name and addressing, another way is to access data by name regardless of the location. Information-centric networking (ICN) evolves the Internet architecture by introducing uniquely named data, so data

is independent from location, application, and means of transportation, enabling in-network caching and replication [5].

The Locator/ID Separation Protocol (LISP) as defined in RFC 6830 (https://tools.ietf.org/html/rfc6830) was published by IETF as an experimental RFC in 2013. LISP separated IP addresses into two numbering spaces: Endpoint Identifiers (EIDs) and Routing Locators (RLOCs). IP packets addressed with EIDs are encapsulated with RLOCs for routing and forwarding in the network. The EID-to-RLOC mapping is stored in a mapping database.

Figure 10.2 shows a typical deployment of LISP in the global Internet [6].

10.2.3 Current Internet Addressing Techniques

The introduction of IPv6 supports of the next generation of wireless, high-bandwidth, multimedia Internet applications, as well as growth in the global number of users and devices. Continued IPv6 deployment provides expanded scale, reduced operational costs by utilizing simpler network models enabling new service and application innovations.

General benefits for IPv6 also include reduction in the deployment of network address translation technologies, increasing address capacity for wireless peer-to-peer (P2P) applications. However, several Internet addressing challenges exist, and new requirements are being introduced based on predicted services and future Internet architecture.

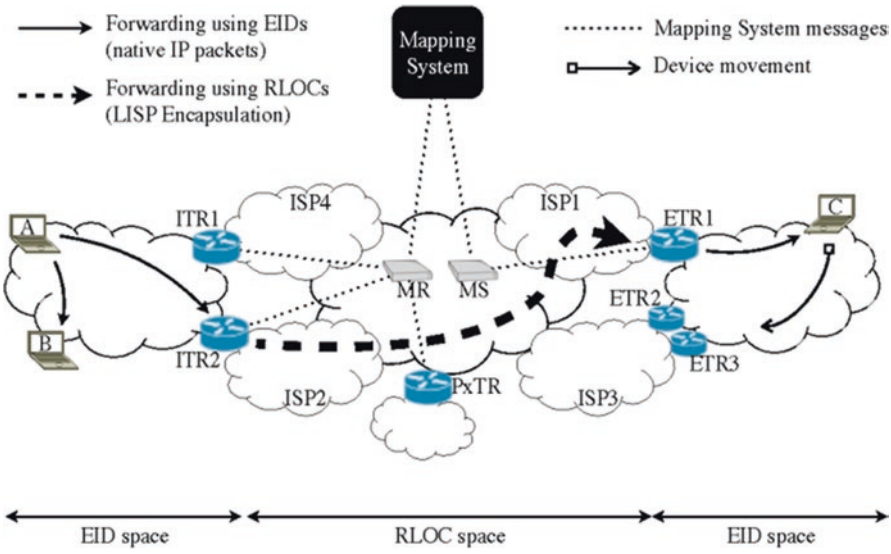

Fig. 10.2 A typical deployment of LISP

10.2.4 Requirements for Addressing in the Future Internet

Several emerging addressing requirements have been identified for future Internet; these include:

- Beaconing for source, destination, and route discovery
- Semantic and contextual addressing
- Flexible addressing
- Inherent security: via identifiers which allow applications and users to validate source and destination

Furthermore, addressing should enable upper layer protocols to identify end points unambiguously, and be agnostic to the underlay technologies and hardware. Thus, it allows the exploitation of new transmission technologies, and decouples users, applications and services from lower-layer hardware. This would also facilitate the interconnection of emerging future Internet devices to existing Internet infrastructure.

10.2.5 Addressing Semantics

Within a limited domain, it is possible to set an address with some special semantic, so service providers or network operators can apply local policies or have certain service bound depending on the sematic. For example, a semantic may denote different device types, or connectivity requirements (https://datatracker.ietf.org/doc/draft-king-irtf-challenges-in-routing/).

The following list shows how address semantics may interact with routing:

- New semantics to IP addresses may have implications for how network routing is performed.
- Semantic techniques might not be supported by existing routing protocols and so would require changes.
- Semantic techniques might enable advanced routing features or offer benefits in scaling and management of routing systems.

10.3 Routing

10.3.1 Network Path Selection

Typically, two approaches may be used for network path selection:

- Firstly, a priori assessment by having the feasible paths and constraints computed in advance
- Secondly, real-time computation in response to changing network conditions

The first scenario may be conducted offline and allows for concurrent or global optimization and several factors to be applied. As network complexity increases, the required computing power may increase exponentially, especially when evaluating large search spaces.

The second approach must consider the speed of calculation. The response processing may delay the service setup, especially if the path selection request is in response to a network failure. Path selection constraints may be applied to reduce complexity. However, the path computation's accuracy and optimality may be negatively affected.

In both scenarios, the amount of information that needs to be imported and processed can become very large (e.g., in large networks, with many possible paths and route metrics), which might impede the scalability of either method.

In the last decade, significant research has been conducted into future Internet architectures. During this research, several techniques emerged, highlighting the benefits of path awareness and path selection for end hosts during this research, and multiple path-aware network architectures have been proposed, including SCION [7] and RINA [8].

When choosing the best paths or topology structures, the following criteria may need to be considered:

- Method a path, or path set, is to be calculated, e.g., a path can be selected automatically by the routing protocol calculated the best path or imposed by a central entity, for example, for traffic-engineering reasons.
- What criteria are used for selecting the best path, e.g., classic route preference, or administrative policies such as economic costs, resilience, and security, and if requested, applying geopolitical considerations.

10.3.2 Traffic Engineering

A fundamental capability of the Internet is to route end-user traffic from the source to the ultimate destination. Transit routers along the path will implement control and optimization techniques to steer traffic along the path from the source to destination. Routers may utilize traffic engineering (TE) techniques to apply scientific principles to the measurement, characterization, modeling, and path selection control and ensure the end-user traffic is forwarded and end-user application requirements are met.

Another objective of Internet TE is to facilitate reliable network operations, which can be achieved by providing mechanisms that enhance network integrity and embrace policies that emphasize network survivability in the event of failures, thus, reducing susceptibility to network outages arising from errors, faults, and physical failures and force majeure events, occurring within the network infrastructure.

Traffic engineering techniques can be applied using distributed or centralized control plane techniques. Both scenarios would utilize the key TE components, including path steering, policy, and resource management:

- Path steering: This is the ability to forward packets using more information than just knowledge of the next hop.
- Policy: This allows for the selection of next hops and paths based on information beyond basic reachability.
- Resource management: This controls how different resources can be shared among different services.

10.3.3 Predictive Routing

Predictive routing means the change in the state of a router/host can be predicted; hence the routing algorithm can make route changes before or as an event occurs. There are new categories of applications that may benefit from predictive routing: such as cars driving on a highway or robots moving in a factory. These are applications where packet loss or delay is potentially very harmful, but the device's movement can be either predefined or predicted.

An alternative approach that alleviates the effects of slow routing protocol convergence is embodied by protocols with packet-carried forwarding state, such as SCION [7] or Segment Routing [9]. In such protocols, forwarding information that is carried in the packet header does not rely on router's (inter-domain) forwarding tables and thus avoiding inconsistent forwarding table state due to asynchronous update mechanisms. Moreover, the nature of the path exploration process in SCION (referred to as beaconing) which creates path segments does not require any convergence for connectivity—instead, additional paths are created over time that become available. Basic end-to-end connectivity, however, is established based on the initial path segments that are disseminated.

In general, the network infrastructure is fixed subject to the impacts of failure, maintenance, and upgrades. However, there is a new class of network emerging based on the use of mobile network infrastructure components such as large constellations of low earth orbiting satellites. These have the property that while the network infrastructure is dynamic, and the best paths are constantly changing, the best path is predictable in advance. This allows a new approach to routing based on current knowledge of the future disposition of the infrastructure rather than on preconfigured "static" paths, or dynamically discovered paths.

10.3.4 ManyNets and Routing for Space-Based Networks

While there is no relation between wireless mesh network routing challenges and protocols developed in IETF MANET WG, routing in space with LEO satellite constellations presents domain specific routing challenges.

A system, where complete global connectivity is provided through LEO satellites, which includes inter-satellite connectivity using Free Space Optical (FSO) transmission, introduces unique set of challenges w.r.t routing in space and possible traffic engineering [10]. This is because (as noted earlier) of the continuous changes to the network paths as the nodes in the orbit are on the move. There is no routing protocol today which does shortest path computation when all the nodes in the network are continuously moving. However, one characteristic of this network is the movements of satellites are completely predictable and this can be factored for new route computation methods. This also introduces unique set of Fast ReRoute (FRR) challenges which are not applicable for terrestrial networks. However, it is worth noting at this time the applicability and possible deployment of such a system is constrained by free space optics (FSO) limitations. These limitations concern with the inter-satellite link capacity, which is currently in the order few Gbps [11, 12], while the sub-sea fiber optical cable provides bandwidth in the order of 10's of Tbps.

The resulting low Earth orbit (LEO) constellations will not only bridge the digital divide by providing service to remote areas, but they also promise much lower latency than terrestrial fiber for long-distance routes. Unlocking this potential is nontrivial: such constellations provide inherently variable connectivity which today's Internet is ill-suited to accommodate. In fact, the use of the BGP protocol to integrate the satellite network in today's Internet unfortunately encounters several major challenges:

- The highly dynamic nature of ground station to satellite links creates.
- Scalability limitations for BGP, especially due to weather disruptions.
- During early phases of deployment, connectivity will fluctuate so often that slow routing convergence with BGP could make the partially deployed constellation unusable.
- The higher cost and lower bandwidth of satellite network links complicates their use for all data traffic, thus complicating the management of differentiated traffic.

There have been proposals to address these challenges. Giuliari et al. propose an optimal solution based on the SCION path-aware-networking architecture, and given this clean-slate baseline, they then develop a more pragmatic solution based on a CDN-like architecture [12].

10.3.5 Mobility

Mobility needs to provide ubiquitous connectivity to mobile users, independent of type and location of devices, access technologies, etc. A mobile node must be able to continue to communicate with others when access location or technology changes when moving and still providing efficient content delivery and trustworthiness.

There have been researches and proposals on mobility for years. One current approach to mobility issues is that they are resolved by the applications themselves using technologies such as MPTCP, QUIC (https://datatracker.ietf.org/doc/html/rfc9000), etc. at the transport layer. Another approach is the Mobile Ad hoc Networks (MANETs) (https://datatracker.ietf.org/wg/manet/about/), which is to provide a network layer solution to support node motions, including IP routing protocol functionality suitable for wireless routing applications.

For the Internet of Everything (IoE), the collaboration of IoE-based devices with current Internet protocols is challenging, specifically in terms of mobility and scalability.

Mobility scenarios in cellular networks pre-REL15 [13] involves only access layer, i.e., UE's mobility from one NodeB to another NodeB with the same or different Mobility Management Entity (MME). However, 3GPP REL15 [13] presents various mobility scenarios which involves IP address changes with or without service continuity as described in various Session and Service Continuity (SSC) modes. In the scenario, where IP address change causes disruption to session continuity, to maintain service continuity, various solutions are specified in [13], involving changes to transport layer protocols at UE. While other category of such solution involves network-assisted service continuity with multiple PDCP sessions and stitching these sessions in backhaul network to prevent the services interruption at the UE without any or with minimal packet loss.

However, there are not widely accepted/deployed solutions in network layer yet for new service requirements described in Gap Analysis of Network 2030" [https://www.itu.int/en/ITU-T/focusgroups/net2030/Documents/Gap_analysis_and_use_cases.pdf]. With the development of new applications in NETWORK2030 with uRLLC requirements, it is desired to support mobility in network layer, which avoids the session interruption and minimizes the packet loss and latency.

10.3.6 Domain-Specific Routing Protocols and Algorithms

New routing protocols are being developed in the IETF for data centers, e.g., RIFT and LSVR. These are protocols specifically optimized for use in certain types of domain and topologies. Such protocols trade general applicability for high performance in the target domain. Soon, there could be more domain-specific cases that require new routing protocols or algorithms, such as routing for satellite communications.

10.3.6.1 Industrial Internet and the Internet of Things

Industrial internet refers to the interconnected networks of sensors, robots, etc. Internet of Things (IoT) network consists of control systems, embedded systems, etc., and in consumer market, IoT is essentially the technology to build smart home and smart cities and enable applications including healthcare, disaster recovery, etc.

New technologies and standards are being developed at a rapid pace to form different IoT ecosystems and networks. From routing perspective, the typical common requirements among these networks are the following:

- Low power consumption. Typical IoT devices are powered by batteries with limited processing power and memory, and this means they need to be conservative on power consumption when sending data packets or control packets. Routing protocols designed for such IoTs should be quiet without sending too many control packets, and then resulted data packets should not have big encapsulation header.
- High availability. Applications such as disaster recovery require the network to provide nondisruptive service in case of network failure, power outage, natural disaster, etc.
- Mobility. IoT devices should be able to connect and communicate with the network or other devices without location and access technology limitations, whenever and wherever.
- Large number of connections. The number of various IoT devices to be connected to the network will be in thousands or millions, so routing protocols are required to connect these huge number of heterogeneous systems.

There are two key issues that future network designers need to contend with in IoT networks. Firstly, the high path quality is needed, which requires the routing system to establish the path and allocate the resources, including the case where it may need to configure the network to strategically replicate and eliminate packets to maximize their chances of successfully traversing the network [14]. Additionally, many IoT devices are designed to meet extreme physical size, cost, and lifetime power budgets. The protocols that these devices use require extreme regard to resource conservation and may not be able to use the "standard" network protocols which are optimized for characteristics such as generality and performance.

10.4 Routing Security and Resilience

Ensuring the security of routing mechanisms continues to be a challenge. Routing attacks include route hijacking, which diverts traffic to an adversary-controlled domain, and denial-of-service attacks, which can prevent communication from happening altogether. Over the past four decades, numerous researchers studied secure

routing in a variety of network types and settings. We briefly highlight the core challenges and several proposed approaches.

An overview of routing security is available as a taxonomy for secure routing protocols by Hollick et al. [14], which emerged from a recent Dagstuhl seminar on secure routing [15]. The taxonomy establishes the following general services that need to be protected: identity service, routing service, topology service, and transport service. An adversary can have a variety of capabilities, resources, and goals—the security section lists different categories of capabilities and resources as defined in Sect. 10.5 of this document. In the context of routing, the main goals are to violate the following security properties: availability of routing and forwarding, authenticity of routing information, confidentiality/privacy of routing and topology information, and anonymity of entities (e.g., mobile users could be located via the routing protocol). In terms of security properties of the forwarded packet data, the routing system should prevent the redirection of traffic flows through entities that intend to eavesdrop or alter packet traffic—if communication is already passing through a malicious entity, it is the responsibility of the data plane to ensure traffic secrecy and integrity.

Routing protocols, especially IGPs, have been running in a relatively benign environment. With the development of new applications, it is critical for the network to provide nondisrupted service especially to high-value traffic. As more hosts/IoTs are added to the network, security is becoming more and more critical.

Secure intra-domain routing protocols have been largely neglected compared to inter-domain settings, as one assumes a benign environment under single administrative control in these settings. In existing intra-domain protocols, however, adversaries can launch several attacks: availability, denial-of-service, or traffic redirection. The typical approach for securing link-state intra-domain routing protocols is to attach a cryptographic signature to link-state updates, as is done for instance in secure OSPF. Within a single administrative domain, the entity identification problem is simplified, as the network administrator can establish and distribute cryptographic keys and certificates among networking devices and systems.

Inter-domain secure routing continues to be a challenge up to today. While S-BGP and its successor BGPSEC have been developed over the past 20 years, they have seen limited deployment due to several reasons: worse scalability than BGP (due to the inability for prefix aggregation and the need for periodic dissemination of routing updates), operational challenges (obtaining and handling certificates, updating router software and possibly even hardware), limited security benefits (new attacks are made possible), slower convergence than BGP, and disruption of policy mechanisms (ASpath alteration/prepending). A beacon of hope is the resource public-key infrastructure (RPKI), which provides the prefix and AS certificates in BGPSEC, as it enables route origin validation, which is easier to deploy than full BGPSEC and in itself addresses several attacks [16]. Unfortunately, the RPKI introduces a circular dependency with routing, as route message verification requires RPKI certificate validation and RPKI certificate validation requires a route to a server to fetch the RPKI certificate database. Moreover, the RPKI also opens up

vulnerabilities to misbehaving RPKI authorities, where a misconfiguration or malicious action can result in rendering an address range unreachable [7].

It appears that an Internet redesign is needed to resolve the thorny issues to secure BGP. The SCION secure Internet architecture has thus redesigned the routing and PKI infrastructure from ground up to achieve high levels of security [7]. By avoiding inter-domain forwarding tables on routers and utilizing a path exploration system that does not rely on convergence, many attacks and vulnerabilities are prevented by design. The control-plane PKI in SCION is constructed such that the distribution of cryptographic credentials follows the transmission of routing messages, thus avoiding circular dependencies between routing and certificate distribution. The definition of trust roots within each isolation domain ensures operational sovereignty and prevents external entities to affect operation due to misconfigurations or misbehavior. As a consequence of its design, SCION can prevent all known routing attacks.

Current routing protocols are built and operated on the assumption of a high degree of trust. IGPs are typically running within a controlled and secured domain and BGP connected with trusted neighbors. For future networks, there are three possible solution directions (not exclusive of each other):

- Making existing routing protocols more secure by adding new authentication mechanisms/algorithms, etc.
- Securing and authenticating the information distributed by routing systems (such as by RPKI mechanisms applied to BGP)
- Using a new secure routing protocol, e.g., SCION

In case of link or node failure, routing protocols should be able to continue to provide an acceptable level of service. This could be achieved through local repair techniques, such as Loop-Free Alternate (LFA) Fast Reroute (FRR) [17, 18]. Meanwhile, routing protocols should reconverge fast and bring the network back to a stable state.

Mutually Agreed Norms for Routing Security (MANRS) (https://www.manrs.org) is a global initiative, supported by the Internet Society, and is made up of network operators to improve global routing security.

MANRS provides critical fixes to reduce the most common routing issues, outlining four simple but concrete actions for network operators:

- Anti-spoofing: Prevent traffic with spoofed source IP addresses. Network operators should enable source address validation and prevent packets with incorrect source IP address from entering and leaving the network.
- Filtering: Prevent propagation of incorrect routing information. Implementing prefix filters within a network can help protect against threats such as prefix hijacking and route leaks.
- Coordination: Facilitate global operational communication and coordination between network operators, maintain globally accessible, and up-to-date contact information.

- Global validation: Facilitate validation of routing information on a global scale. Network operators need to ensure that their network's routing information is publicly available including the announcements that the network originates and the routing policy.

Figure 10.3 shows a summary of proposed actions by MANRS (https://www.manrs.org)

10.5 Emerging Routing Protocols

Internet paths often require evaluating and assessing route metrics, including latency, jitter reliability, bandwidth, and congestion. Depending on the number and overall path length, computing paths is often processor and time-consuming. The design of effective path evaluation strategies is a balancing act between accuracy, computation time, and be more specific.

Several emerging routing techniques are being developed to address the scaling concerns and path selection complexity for the future Internet and support emerging domain-specific technologies. These new routing techniques are discussed in the following subsections.

A CLOS network [19] is a kind of multistage circuit-switching network. It was invented by Charles Clos to solve the problem of explosive growth of telephone network. Figure 10.4 shows a Clos network, where each leaf is connected to every spine node, and vice versa.

Modern data centers, especially large-scale data centers, host tens of thousands of end points, and this puts on new challenges on network architecture and routing protocols. For operational simplicity, many data centers have chosen BGP [2] with a CLOS topology as the most appropriate routing protocol and architecture as described in RFC 7938 (https://tools.ietf.org/html/rfc7938).

Fig. 10.3 Proposed actions for service providers by MANRS

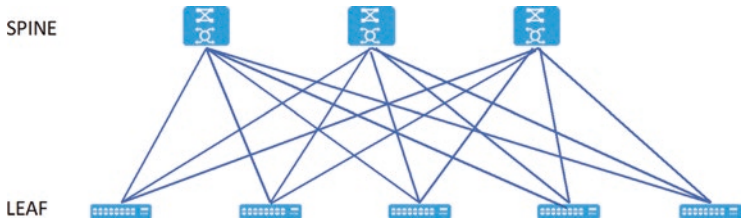

Fig. 10.4 A CLOS network architecture

10.5.1 RIFT

RIFT (Routing In Fat Trees) is a novel routing protocol defined by IETF (https://datatracker.ietf.org/wg/rift/about/). It mainly targets Clos [19] and fat-tree network topology-based data centers and is optimized with minimization of configuration and operational complexity.

RIFT is mixture of both link-state and distance-vector technologies and can be described as "link-state towards the spine" and "distance vector towards the leaves."

Here are the major characteristics of RIFT:

- Northbound link state routing with flooding reduction, lower levels are flooding their link-state information in the "northern" direction, so that each level obtains the full topology of levels south of it.
- Southbound distance vector routing, each upper node generated a default route to the "southern" direction.
- Link state is advertised one-hop southbound and then reflected one-hop northbound. This is when a node detects that default route encompasses prefixes for which one of the other nodes in its level has no possible next-hops in the level below, and it has to disaggregate it to prevent black-holing or suboptimal routing through such nodes.
- Optional zero touch provisioning (ZTP), only top tier nodes need to be configured.
- Packet formats are defined in Thrift [20] models.

Figure 10.5 is a simplified illustration of the RIFT protocol (https://datatracker.ietf.org/meeting/103/materials/slides-103-rtgarea-rift-update).

10.5.2 LSVR

The link state vector routing (LSVR) working group (https://datatracker.ietf.org/meeting/103/materials/slides-103-rtgarea-lsvr-update) at IETF is proposing a new solution which leverages BGP link-state distribution and the Shortest Path First (SPF) algorithm and targets Massively Scaled Data Centers (MSDCs).

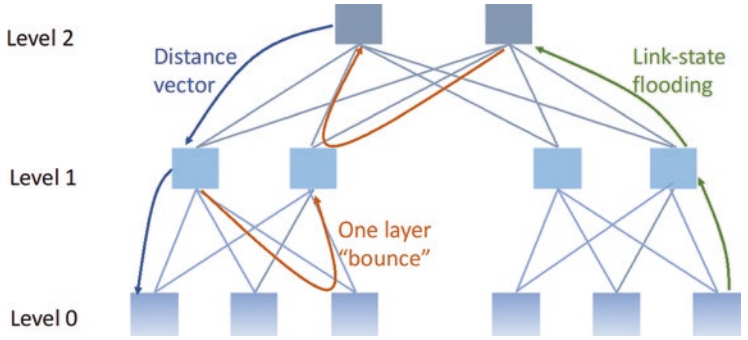

Fig. 10.5 Illustration of RIFT: routing in fat trees

BGP has been chosen as the single routing protocol to simplify routing in MSDCs and their interconnections. While BGP offers operational simplicity and scalability, it lacks some advantages that IGP can provide, such as being a hop-by-hop routing protocol, and it is missing the fabric topology information as IGP. As the size of data center grows, the CLOS tiers increase as well as configuration complexity.

Link state vector routing, also known as BGP-SPF, using BGP as base protocol, and add the best of IGP characteristics, and the following are the main advantages:

- Complete fabric topology at each node, and path computations using SPF, TE, CSPF, LFA, etc.
- Faster convergence compared with classic BGP
- Simplicity—incremental from base BGP, operational and troubleshooting
- Incremental updates, no flooding and selective filtering
- Reliable transport using TCP

The key idea of BGP-SPF is that BGP runs Dijkstra algorithm to calculate best path instead of the BGP Bestpath decision process. BGP-SPF supports various peering models, such as peering in single-hop or route-reflector, as long as all BGP speakers in the BGP-SPF domain can receive link-state NLRI and hence perform consistent distributed route computing. To be backward compatible, BGP-SPF introduces the BGP-LS-SPF SAFI for BGP-LS SPF operation and extends BGP-LS (https://tools.ietf.org/html/rfc7752) with new attribute TLVs.

Figure 10.6 shows the building blocks of LSVR protocol:

10.5.3 SCION

The SCION (Scalability, Control, and Isolation On Next-generation networks) inter-domain network architecture has been designed to address security and scalability issues and provides an alternative to today's BGP. SCION combines a globally distributed public key infrastructure, a way to efficiently derive symmetric keys

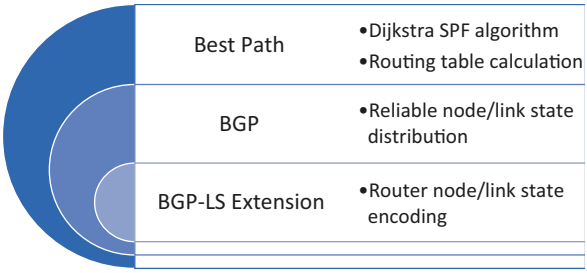

Fig. 10.6 Components of LSVR

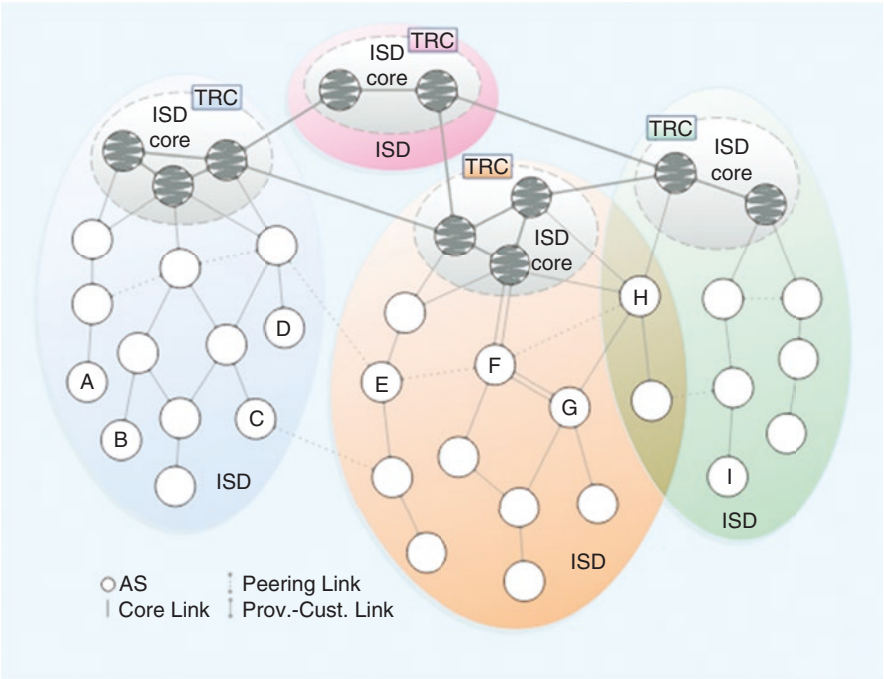

Fig. 10.7 SCION architecture overview

between any network entities, and the forwarding approach of packet-carried forwarding state. Instead of relying on inter-domain lookup tables, the AS-level forwarding path is encoded in the header of the packet. Each router verifies a message authentication code with a symmetric cryptographic key before forwarding. Figure 10.7 depicts the ISolation Domains (ISDs), which group a set of ASes into independent routing domains. The ISD is governed by a set of core ASes that provide connectivity to other ISDs, a role that is typically held by the largest ISPs that also provide global connectivity in today's Internet. The Trust Root Configuration (TRC) of each ISD enables setting of the local roots of trust, in essence the set of public keys that are used to verify public-key certificates. The partition into ISDs also enhances the scalability of the SCION routing system, as the beaconing mechanism

creates intra-ISD and inter-ISD path segments among a reduced number of entities compared to BGP.

The SCION Internet architecture provides a fundamentally clean-slate approach to multipath communication: at the control plane, the routing system discovers a variety of AS-level path segments (which can also differ in the interface or links connecting neighboring ASes), which are globally disseminated through a path server infrastructure; at the data plane, cryptographically protected packet-carried state encodes the AS sequence and the AS-to-AS interfaces in the packet header.

End-hosts fetch path segments from the path server infrastructure and construct the exact forwarding route themselves by combining those path segments. The architecture ensures that a variety of combinations among the path segments are feasible, while cryptographic protections prevent unauthorized combinations or path segment alteration. The architecture further enables path validation, providing per-packet verifiable guarantees on the path traversed.

SCION's intrinsic multipath communication provides a natural defense against distributed denial of service (DDoS) attacks. An attacker must congest all paths instead of only one, which increases the needed attack capacity and complicates the attack since access to all paths must be prevented. Further, an AS can choose not to publicly announce some of its path segments at the path servers, but still share them with select communication partners "out of band." The ability to use such "hidden" path segments as part of multipath communication guarantees the existence of a fall-back path that is not publicly known and therefore cannot be clogged through a DDoS attack.

10.6 Conclusions

A set of requirements and several major innovations have been outlined in this chapter for future Internet routing and addressing. Fundamentally, the recurring theme for future routing and addressing is the adherence to a key set of principles which must be preserved and applied to the future architecture of the Internet, and these include:

- Heterogeneity support

  Given existing knowledge of Internet evolution, we must assume the requirement for heterogeneity to be much higher than it is today. Multiple types of devices and applications, network nodes, and protocols will coexist. Hence, the capability to support heterogeneity should remain and potentially be an enforced requirement.

- Massive scale, throughput, and scalability

  With the deployments of massive IoT devices and large-scale data centers, the number of devices needs to be supported by routing protocols that keep increasing. Protocol design and extension have to put scalability into consideration.

New applications and services also add new challenges, such as application specific service requirements in terms of latency and throughput. Routing protocols should have the capability to provide diverse routing services to meet the needs.

- Autonomic networking

  Generally, device and connection management require a series of manual processes and expert knowledge, although more recently the advent of device and service models, along with well-defined APIs, is facilitating the rise of network automation.

  Future services and increasing need for traffic-engineering will require the network to be more flexible and adaptive. As new algorithms and protocols are then proposed, it will also require additional configuration steps and models. Future Internet network devices will need to be adaptive but also more autonomic, with the capability to auto-configure and self-heal.

- Security and resilience

  Internet has changed people's everyday life dramatically, and this also means an ever-increasing dependency of the Internet. To provide reliable and secured services is key requirement for networks. Routing security, as an essential piece of a secured Internet, continues to be a challenge. Network resilience is to maintain an acceptable level of services against failure or damage, and for routing protocols, this means to reconverge fast when failure happens on top of various local repair techniques.

Fundamentally, future Internet must also stay true to the principles that have made the existing Internet so successful, which include openness and decentralization.

References

1. Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473). International Standard 10589: 2002, 2nd ed, 2002
2. Y. Rekhter, T. Li, S. Hares, A border gateway protocol 4 (BGP-4). RFC 4271 (2006), https://www.rfc-editor.org/info/rfc4271
3. *A History of the ARPANET: The First Decade (Report)* (Bolt, Beranek & Newman Inc., Arlington, VA, 1981)
4. Internet Protocol, Version 6 (IPv6) specification, RFC8200, TBC
5. IRTF Information-Centric Networking Research Group (ICNRG)
6. D. Saucez, L. Iannone, O. Bonaventure, D. Farinacci, Designing a deployable internet: the locator/identifier separation protocol. IEEE Internet Computing Magazine **16**(6), 14–21 (2012)
7. A. Perrig, P. Szalachowski, R.M. Reischuk, L. Chuat, *SCION: A Secure Internet Architecture* (Springer, New York, 2017)
8. J. Day, *Patterns in Network Architecture: A Return to Fundamentals* (Prentice Hall, Hoboken, 2008)

9. C. Filsfils, S. Previdi, L. Ginsberg, B. Decraene, S. Litkowski, R. Shakir, Segment routing architecture, RFC 8402 (2018), https://www.rfc-editor.org/info/rfc8402
10. D. King, A. Farrel, Z. Chen, An evolution of optical network control: from earth to space, in *22nd International Conference on Transparent Optical Networks (ICTON)*, (ICTON, Bari, 2020), pp. 1–4. https://doi.org/10.1109/ICTON51198.2020.9203098
11. P. Miller, Ka-Band – the future of satellite communication, http://www.tele-satellite.com/TELE-satellite-0709/eng/feature.pdf
12. Giacomo Giuliari, Tobias Klenze, Markus Legner, David Basin, Adrian Perrig and Ankit Singla. Internet backbones in space. In ACM SIGCOMM Computer Communications Review ACM New York, 50(1), 2020.
13. TS 23.501 System architecture for the 5G System (5GS)
14. A. Herzberg, M. Hollick, A. Perrig, Assessment of the effective performance of DPSK vs. OOK in satellite-based optical communications, ICSO 2018. Dagstuhl Rep. **5**(3), 28–40 (2015). https://doi.org/10.4230/DagRep.5.3.28
15. D. Cooper, E. Heilman, K. Brogle, L. Reyzin, S. Goldberg, On the risk of misbehaving RPKI Authorities, in *Proceedings of ACM HotNets-XII*, (ACM, New York, 2013)
16. R. Lychev, S. Goldberg, M. Schapira, Is the juice worth the squeeze? BGP security in partial deployment, in *Proceedings of ACM SIGCOMM* (2013)
17. M. Shand, S. Bryant, A framework for loop-free convergence, RFC 5715 (2010), https://www.rfc-editor.org/info/rfc5715
18. Topology independent fast reroute using segment routing, https://tools.ietf.org/html/
19. C. Clos, A study of non-blocking switching networks. Bell Syst. Tech. J. **32**(2), 406–424 (1953). https://doi.org/10.1002/j.1538-7305.1953.tb01433.x
20. Apache Software Foundation, Thrift interface description language, https://thrift.apache.org/docs/idl