7

DHCP SERVER DEPLOYMENT STRATEGIES

This chapter examines deployment strategies and trade-offs for DHCP servers. Most trade-offs pit budget dollars against quantities of servers, so the most common goal is to deploy DHCP servers where end users will always be able to obtain these services in a timely manner, while minimizing the total dollars spent on servers deployed and associated server lifecycle expenses. This simply stated goal implies a need for highly available and reasonably performing services, all provided within budget constraints. Budget dollars must account not only for server purchases, but for ongoing support and maintenance, which includes server hardware upgrades, operating system (OS) patches and upgrades, as well as DHCP upgrades for new features, bug fixes, or security measures.

7.1 DHCP SERVER PLATFORMS

DHCP servers can be deployed in a variety of platforms from physical hardware servers or appliances, or as virtual servers on a virtual machine (VM) platform. When we discuss deployment options, we'll generically use the term platform, which can generally be interpreted as either one of these options in each case.

IP Address Management: Principles and Practice, by Timothy Rooney

7.1.1 DHCP Software

The traditional model for deploying DHCP servers entails deploying a physical server supporting the recommended processing components and operating systems supported by the corresponding DHCP software vendor. Other applications may also be installed on such servers to maximize hardware utilization.

7.1.2 Virtual Machine DHCP Deployment

Virtual machines (VMs) exist for major Windows and Linux operating systems (OSs), enabling the deployment of Microsoft DHCP on Microsoft VMs and ISC on Linux VMs. Deployment on VMs saves on hardware costs, rack space and power draw, while enabling better segregation than in installing a DHCP daemon on a generic hardware server. Major appliance vendors also offer their appliance products as virtual machines, combining the benefits of VMs with the benefits of appliances, discussed next.

7.1.3 DHCP Appliances

DHCP appliances are preinstalled DHCP services on secure hardware platforms, typically Intel-based platforms with a hardened Linux operating system. Like routers, which were initially deployed as software running on general purpose hardware and evolved to special purpose hardware platforms, DHCP appliances offer an evolutionary path to self-contained hardware platforms for DHCP services. Appliances are "hardened" in that the base Linux kernel installed on the platform has been stripped of any unnecessary services. This results in a customized kernel and OS that supports only DHCP services (and other services supported by the vendor such as DNS). Underlying file system, users, permissions, and network ports should also be pared down accordingly by the appliance vendor.

Appliances offer simplified deployment with one-stop shopping, instead of having to coordinate and acquire server hardware, install the proper OS version and patch levels, then install DHCP services software. Appliances can simplify the ongoing upgrade process by prepackaging upgrades with compliant OS and services versions with corresponding hardware platforms. Depending on the vendor, these upgrades may be applied from a single centralized console, eliminating the need to physically deploy staff to perform upgrades. In addition, most vendors support centralized monitoring of deployed appliances, enabling proactive detection of outages or degradations.

Of course appliances generally cost more than general purpose server hardware, and most incorporate ISC DHCP services, which are freely available for most leading OSs at www.isc.org. In this chapter, we'll focus on deployment strategies for DHCP services, regardless of implementation on general hardware or appliances platforms.

7.2 CENTRALIZED DHCP SERVER DEPLOYMENT

The deployment of DHCP servers generally comes down to a trade-off between wide distribution of a large number of servers "closer" to clients versus narrow distribution of

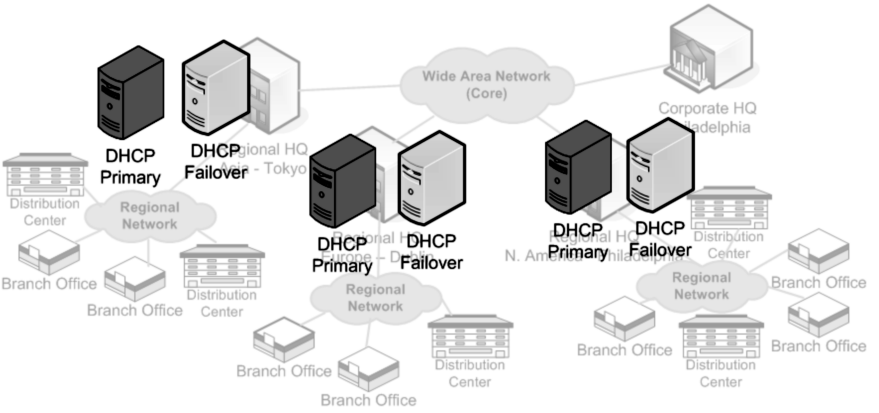

Figure 7.1. Centralized DHCP server deployment for IPAM Worldwide.

a fewer number of DHCP servers serving clients from a variety of locations. The extremes of this trade-off consist of a DHCP server on every subnet versus one or more DHCP servers centrally located serving all of the organization's clients. The key is to balance availability and reasonable performance of the DHCP service between clients and servers while remaining within budget constraints for servers and ongoing management thereof. Your deployment will likely fall between these two extremes.

Figure 7.1 illustrates the fully centralized deployment approach scenario for IPAM Worldwide. Overlaying the high-level network diagram from Figure 3.2 in Chapter 3, this scenario features the deployment of a pair of DHCP servers per region, one functioning as the primary and the other as failover or backup. All DHCP traffic must be funneled to the regional headquarters sites, imposing higher reliance on robust network connectivity to these sites from the respective regions. This architecture also implies the DHCP server hardware is sufficiently sized to meet performance and capacity requirements. Note that DHCP primary and failover servers should generally be deployed in separate physical locations for disaster resilience. An outage at one site would not interrupt all DHCP services for a region.

7.3 DISTRIBUTED DHCP SERVER DEPLOYMENT

At the other end of the deployment continuum, the decentralized deployment approach is illustrated in Figure 7.2. In this figure, a primary DHCP server is located at [nearly] every branch office and distribution center. This localizes DHCP traffic, affording deployment of less stringently sized DHCP servers. Network connectivity to the regional headquarters is still required however due to the deployment of DHCP failover servers there. These servers act as failover servers for the regional servers, though more than one per region may be required for load sharing. Consider the load and redundancy capabilities of your chosen DHCP vendor to identify viable alternative architectures for your network.

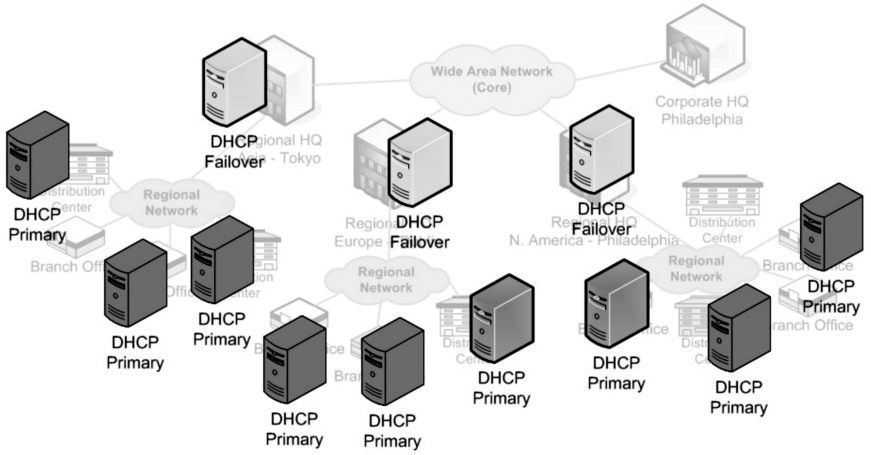

Figure 7.2. Distributed DHCP server deployment for IPAM Worldwide.

Contrasting the two extremes of Figures 7.1 and 7.2, the former requires fewer, albeit more powerful DHCP servers and rock solid network connectivity to the regional headquarters sites. The latter requires many more DHCP servers, though of more modest specifications, providing localized services with a network reachable shared backup. You may be wondering, if the network link to a site goes down, what good is having an IP address from a DHCP server? Without a redundant link, other than providing IP access to local network resources, it may indeed be of limited value. But in the centralized architecture, without distributed sites, if a link to a regional headquarters site fails, clients requiring new or renewed address leases will likewise be rendered useless. As always, the trade-off must be considered and generally a mixed approach of centralized with at least partial distribution often minimizes overall outage risk.

While the ISC DHCP server is a single-threaded application, its performance is usually sufficient for most environments. If you have several thousand DHCP clients attempting to obtain leases at about the same time however, some delays will be likely. If this occurs frequently, you may want to consider deploying additional servers and partitioning finer networks-per-server granularity to reduce the load per server. Again, this is usually not a major concern unless you are a service provider utilizing DHCP to initialize devices like customer premises modems for paying subscribers. After recovery from a neighborhood power outage, devices will come back up and inundate the DHCP server for addresses. In such environments, it probably makes sense to consider a commercial performance-oriented DHCP server.

Prepare your routers to support DHCP by configuring the IP addresses of your DHCP servers within your routers' relay agent lists. These lists within each router enable the router to terminate received DHCPDISCOVER packet broadcasts, then retransmit them as unicast packets to each configured DHCP sever IP address on its relay agent list. If you partition your network such that address pools for certain subnets are served by a given DHCP server, while those for other subnets are served by another DHCP server, make

sure you configure routers serving those subnets accordingly. You could add all DHCP servers to all routers, but this will result in needless relay agent traffic, especially if you have several DHCP servers. DHCP for IPv6 networks utilizes well know multicast addresses, obviating the need to configure relay agent lists on routers, though such configuration may alternatively be performed on the relay agent to control which DHCPv6 servers are to process relayed DHCP transactions and not just any DHCPv6 server listening on this multicast address.

7.4 SERVER DEPLOYMENT DESIGN CONSIDERATIONS

Key considerations when formulating the DHCP server deployment design including the following:

- *Response Time Requirements*. Do your clients have stringent response time requirements? Most popular clients tolerate response times in the seconds, but certain applications may be more demanding. The more stringent your requirements, the more important will be server performance and perhaps client proximity.
- *Load Requirements*. Do you have certain load conditions that must be handled? For broadband service providers utilizing DHCP as a customer premises equipment initialization technology, load spikes may occur upon recovery from a residential power outage or equipment installation or reboot. For enterprise environments, such a spike could occur at the start of the workday if several associates arrive at or near the same time, though many devices will simply attempt to renew an IP address previously used by default.
- *Traffic Expectations*. Do you employ short lease times to minimize overbooking, which causes more frequent renewal attempts? Generally the shorter the lease time (the T1 and T2 times), the shorter the interval between obtaining the lease and subsequent lease renewal attempts. This drives increasing traffic on the network to and from the DHCP server(s) and must be considered when designing to the aforementioned response time and load requirements for server quantities and associated bandwidth.
- *Availability Requirements*. Do your clients positively have to be able to obtain an IP address or configuration via DHCP 24×7 or is the service "best effort" based? Most will answer that high availability is critical, but with devices growing increasingly multinetworked, as long as one network's address assignment mechanism is available this may be acceptable.* Mean time to repair (MTTR) is another consideration in meeting DHCP services availability objectives. Having a spare server locally can shorten MTTR while having to order a replacement will delay this process.

* Of course this statement assumes different DHCP services serve these different interfaces that may not be the case.

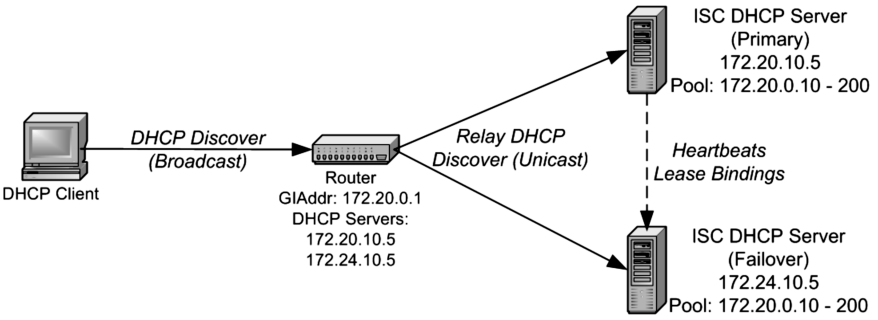

Figure 7.3. ISC DHCP failover configuration.

The first three considerations above relate to deploying sufficient quantities of servers of given lease distribution rate to meet respective performance objectives. A good starting point is to identify the number of expected DHCP clients at each site on your network. This number should account for all devices requiring DHCP, including data devices, voice devices, and all IP devices requiring DHCP at each site. Don't forget to account for "peak" quantities of users and devices so that everyone, even associates visiting on temporary basis, may obtain a valid lease.

After accounting for peak quantities of DHCP clients, consider the frequency of DHCP transactions. This will be dependent on your lease times, as well as client lease release configuration. Most clients will "remember" a prior lease and attempt to request it upon power-up, for example, when an employee returns to work the next day, though this is not always the case.

The fourth consideration listed above relates to providing high availability DHCP services for DHCP clients. Given the general importance of providing highly available DHCP service, deploying for high availability is typically recommended. Once you've designed your deployment based on performance requirements, total or selective high availability may be planned. Based on server technologies you plan to deploy, implementation of high availability will impact not only the number of servers required, but potentially your address space plan.

The ISC and the Microsoft DHCP implementations utilize vastly different approaches. The ISC server employs a failover protocol* such that for a given address pool, one DHCP server will act as the primary, while a second DHCP server will act as the backup or failover server. This basic configuration is illustrated in Figure 7.3.

Each relay agent must be configured to unicast received DHCP [for IPv4] broadcast packets to both the primary and failover DHCP servers, 172.20.10.1 and 172.24.10.1 in Figure 7.3. Recall that DHCPv6 relay agents can likewise be configured with DHCPv6 server addresses or may utilize a well-known site-scoped multicast address, FF05::1:3. The DHCP servers utilize a failover protocol such that the primary sends heartbeat

*The ISC implementation was based on RFC draft specifications by the IETF, which were largely tabled. However, the IETF is endeavoring to redefine the DHCP failover protocol and ISC plans to implement the new version while also supporting the current RFC draft-based implementation.

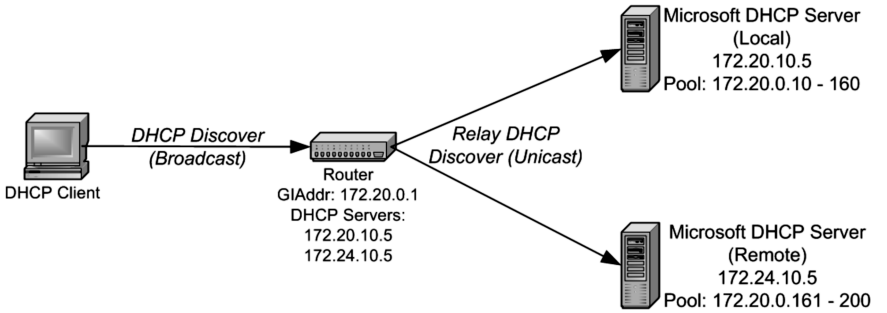

Figure 7.4. Split scopes configuration.

messages as well as lease binding information to the failover server. The failover server utilizes user-settable parameters to determine that the primary is down and begins processing the unicast DHCP packets from the relay agent(s). Thus, clients are able to continue receiving IP address and parameter assignments despite the primary server being down. Upon recovery, the primary server obtains the current lease database from the failover server, then assumes its role as primary once again.

The Microsoft approach does not employ an interserver protocol like DHCP failover. Instead, by deploying two DHCP servers with complementary address pools, *not the same address pools*, either server can process DHCP transactions without worry of duplicate assignment. Microsoft recommends using the "80-20 rule" in configuring 80% of the address pool on a "local" server and 20% on a "remote" server. In this way, most of the DHCP transactions will be handled by the local server assuming the client will receive the first offer from the local server and accept it. This configuration is illustrated in Figure 7.4, where we have split our 172.20.0.10-200 pool using the 80/20 guideline.

Like the ISC failover configuration in Figure 7.3, each relay agent needs to be configured with both local and remote Microsoft DHCP server addresses. Each DHCP server is configured with an address pool for the relevant subnet, 172.20.0.0, but with nonoverlapping contents. In this example, we've used the same total pool size for the pair of Microsoft DHCP servers as was used in the ISC example in Figure 7.3. Since both servers are required to meet the capacity needs, you may end up with an inability to meet IP address demands should one fail. Another alternative is to configure the local DHCP server with 100% of the required local capacity and allow overflow of extra addresses to the remote server for backup. In this manner, the local server can handle 100% of the capacity, and the remote can assist with a portion of those additional clients when the local server is unavailable. Referring to Figure 7.4, the local server could be configured with address pool 172.20.0.10-200 and the remote server with 172.20.0.201-254. This additional capacity can range up to 100% of the required capacity, providing 100% redundancy at the expense of doubling the required address space. While popularized by Microsoft, the split scope approach may be used with various vendor DHCP servers.

From a security perspective, implementations of DHCP authentication have not been broadly commercialized. Thus, little practical security exists for securing the

DHCP transaction itself. This may not be of major concern for enterprise networks, where DHCP is provided for internal use, but this may be problematic if users unknowingly start a DHCP service on their machines. Most users probably don't have a DHCP server installed but those with [self-perceived] IT expertise may.

For service provider networks using DHCP to initialize customer premises equipment, the use of the service provider gateway or edge device can provide some assurance as to the validity of the DHCP client for address assignment. Tying DHCP into the provisioning process can help correlate a DHCP client with a paying subscriber identifier to minimize theft of service.

DHCP itself can be used as a means to "secure" network access by determining whether a given DHCP client meets acceptance criteria for admission to the network by virtue of IP address assignment by the server. This provides a form of network access control, though it does not protect against IP address spoofers. Configuring DHCP for access control security is discussed in the next chapter.

7.5 DHCP DEPLOYMENT ON EDGE DEVICES

Most router vendors provide a DHCP service as a component of their router platforms. This may lead one to question whether a separate server is needed to support DHCP services. As with most design questions, the answer is, "it depends." Small environments with a few sites with local routers serving up to 100 or so monolithic clients each may be well served by configuring the router to provide DHCP services. However, larger organizations or those requiring more advanced DHCP services, for example, for discriminating voice versus data clients for address and option parameter assignment, would be better served deploying discrete (nonrouter-integrated) DHCP servers.

The advantages of running DHCP on a router device include the following:

- *Lower Hardware Cost*. No need to procure a server or set of servers.
- *Single User Interface*. The same command line interface can be used to configure the router and the DHCP server, and no relay agent configuration is needed.
- "*Fewer Moving Parts*." One less communication link and server required to perform DHCP functions, which in general can increase the overall solution reliability.

The main disadvantages of running DHCP on a router are as follows:

- *Options Support*. Most router-based DHCP servers are primitive, supporting address assignment but little in the way of options support.
- *Client Class Support*. Major vendors do not support client classes, which are required for discriminatory address/option assignment to different devices, for example, VoIP versus data devices.

- *No Failover*. If a router fails, you've probably lost connectivity in any case but if there are two routers serving a subnet for redundancy a split scopes approach would have to be employed, increasing management complexity.
- *No Centralized Management*. Router-based DHCP services are configured via command line and unless a centralized tool is employed, each router DHCP server must be configured manually with respect to the IP addressing plan; less likely support is possible if multiple router vendor products are in use.