CYGNA LABS

Automating Cloud DHCP-DNS-IPAM (DDI)

By Timothy Rooney

About Cygna Labs

Cygna Labs is a software developer and one of the top three global DDI vendors. Many Fortune 100 customers rely on Cygna Labs' DDI products and services, in addition to its industry-leading security and compliance solutions to detect and proactively mitigate data security threats, affordably pass compliance audits, and increase the productivity of their IT departments. For more information, visit https://cygnalabs.com.

Introduction

The cloud is transforming how organizations offer and support computing and communications services to their users and constituents. Cloud computing enables IT or service provider operations teams to provide network, computing and services capacity with flexibility, efficiency, and elasticity. These benefits are realized primarily through the cloud's characteristic use of virtualization technologies, which enable them to rapidly instantiate additional capacity for a required service element within minutes. If capacity requirements fluctuate over time, capacity can just as quickly be withdrawn or allocated elsewhere. This elasticity affords organizations agility and cost efficiencies in offering network, computing, and services resources dynamically sized to changing capacity needs over time.

Enterprises can leverage public cloud services such as those offered by Amazon Web Services (AWS), Microsoft Azure, Oracle Cloud, Google Cloud, and others. They can also build cloud functionality within their own data center infrastructure. This "private cloud" configuration can be paired with one or more public cloud services in a hybrid implementation that enables the enterprise to support typical busy time services capacity while relying on the public cloud to support particular services or capacity overflow.

Service providers, inherently cloud providers by virtue of supporting networks and services external to their customers, can leverage cloud technologies not only for efficient services capacity management but also for rapid service deployment. Historically, service providers needed to acquire specialized hardware and software to perform functions required for new service offerings. Such capital-intensive investments were necessary to meet service providers' rigorous reliability, availability, and scalability requirements, among others. By virtualizing networkfunctions (i.e., network functions virtualization, NFV), service providers seek to lower capital costs and accelerate time to market by focusing any specialized development to the software function only, while leveraging common robust hardware. Such software defined networks (SDN) offer rapid, cost-effective, and flexible services delivery.

This white paper discusses the foundational elements required for every cloud infrastructure, namely the management of IP addresses and domain name system (DNS) elements in a dynamic cloud environment. These core network services facilitate the initialization of critical network layer configuration which when automated, are indispensable to achieving rapid, scalable elasticity.

Core network services

Core network services are named as such given they are network services in that they service network and computing elements in initializing their respective IP configurations to enable communications on the IP network and core in the sense that without which, such IP communications would be impossible.

This IP configuration information consists of a unique IP address, generally a host domain name, typically as a human-consumable reference to the element, and perhaps additional initialization information such as a bootfile name and location. Thus, each virtual element, whether a network, computing or service element, needs to be provisioned with at least one IP address, hostname and optionally other configuration parameters.

IP Initialization Process

Virtual network functions (VNFs) or generically, virtual machines (VMs), require provisioning of basic IP network information upon instantiation as would any network device upon deployment on the network. As such, core network services are critical for virtual environments. Certainly, each VM requiring network connectivity will require assignment of an IP address or in some instances, multiple IP addresses. In most cases, the VM will be assigned a hostname such that it can be referenced by name instead of its IP address.

This name reference is necessary not only to simplify navigation by humans as mentioned earlier, being able to connect to a VM using its name, but also potentially by other VMs as necessitated by service chaining. For example, a series of VNFs or VMs may be required to provide a given service; if a predecessor VM references its next VM in the chain or series by name, this enables the provisioning of elastic capacity for the function provided by that VM.

Consider the simple example in Figure 1. A given service requires in-band processing through two service components. The service provider[1] can manage the service capacity dynamically through the deployment of these service elements as VNFs/VMs within a private, public or hybrid cloud deployment. A given user's data path may traverse any individual element as long as both functions are traversed. As demand grows or even spikes at a given time, dynamic instantiation of supplemental VNFs/VMs enables deployment of additional service element capacity. When demand subsides, VNFs/VMs may then conversely be decommissioned.

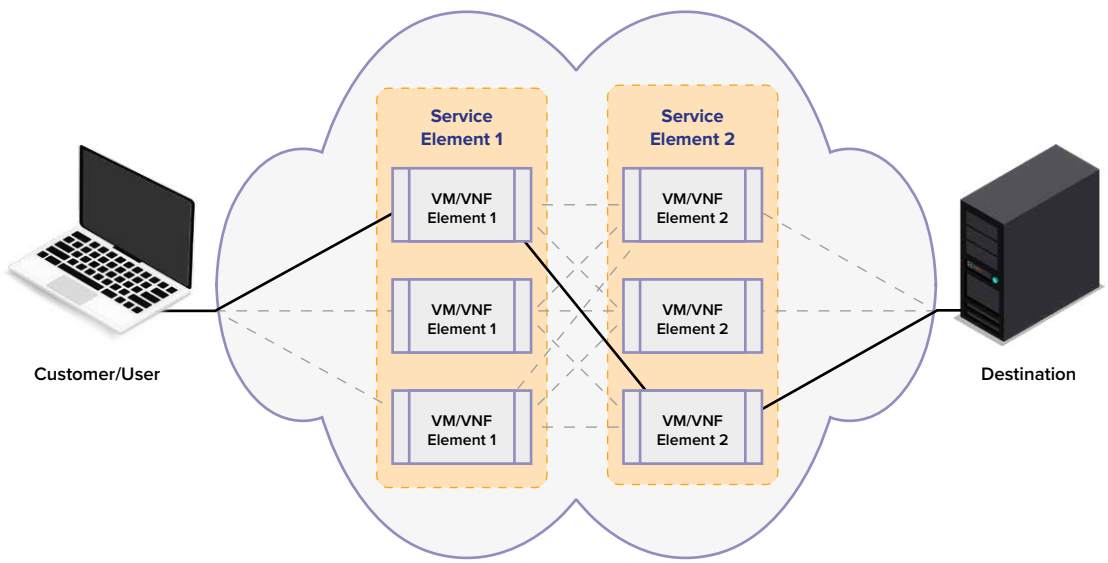

Figure 1: Example service serialization

[1] I'm using the term "service provider" in a broad sense to include IT organizations providing networking services to their constituents.

In providing a service like the one in Figure 1, IP packets from each service element 1 VM must be routed to a service element 2 VM. Instead of configuring each assigned IP address of each downstream (service element 2) VM into each precedent (service element 1) VM, which adds effort, time and the opportunity for error in the provisioning or decommissioning process, one can merely add or remove the created or destroyed VM's IP address from a DNS entry for the given function. Thus a "function" entry in DNS such as service-element-2.example.com can resolve to one, two, or any number of IP addresses, i.e., one per actively provisioned VM. This set of IP addresses associated with a common domain name comprises what's referred to as a resource record set.

As capacity fluctuates based on demand, correspondingly active VM IP addresses can simply be updated in DNS within the function's resource record set. For each new customer flow, the VM need only query DNS to identify a successive node in the path. Of course if a deterministic flow among VMs is required for the service, each VM's unique host domain name and IP address association may be published in DNS instead.

IP Initialization Implementation

Let's consider some approaches to configuring these core network parameters during the process of instantiating a VNF or VM. Among other parameters, each VM, technically each VM interface, requires assignment of a unique IP address, unique at least within the routing domain. When using a public cloud service, public IP addresses are assigned by the cloud provider, while a public cloud customer may allocate a portion of their internal (typically private) address space for use for public cloud infrastructure provisioned as an extension of the enterprise network. However, even when using internal address space within a public cloud, IP address assignments are generally performed by the public cloud platform; hence the tracking of public cloud IP address space requires querying of the corresponding platform.

While public cloud services generally control IP address assignments within the space allocated by the subscribing organization, IP address assignments within a private cloud are generally more tractable. These assignments may be defined manually, on an ad hoc basis as each VM is created, or dynamically leveraging the Dynamic Host Configuration Protocol, DHCP. A DHCP server needs to have been pre-provisioned with an IP address pool, from which individual IP addresses can be assigned upon request.

The benefit of using DHCP is that a pool of IP addresses can be pre-provisioned and assigned to VM's on the fly as capacity needs dictate. Another benefit is that additional IP initialization parameters, such as a bootfile location, can be communicated to the requesting device via DHCP options included with the assigned IP address. The downside is that the DHCP pools must be monitored to assure IP address availability upon demand; if an address pool depletes of its IP addresses, additionally created VMs will be unable to obtain an IP address and successfully initialize for network communications. If DHCP is not used, then a means to assign a static IP address to each VM upon instantiation is required. This may necessitate a manual lookup in a spreadsheet or an automated link to an IP address management (IPAM) system as we'll discuss later.

Using either method for IP address assignment, DHCP or static, an entry in DNS is also likely needed for reasons discussed previously. Updating DNS requires entry of the assigned IP address in association with the VM hostname or function name within the DNS configuration. Such entry can be performed manually by editing the appropriate DNS zone file or by performing an update via a utility like nsupdate. Alternatively, if DHCP is used to assign the VM IP address it can automatically update DNS, though mapping the IP address to the individual VM host name, not to a broader service name.

DHCP Method

Let's consider the process flow for initializing the IP configuration using DHCP as outlined in Figure 2. Along the top of the figure, our three basic macro-level steps entail preparation to instantiate the VM, followed by the action of instantiation, then by completion of instantiation with a unique valid IP address and DNS entry. Supporting this process, as part of the preparation phase, we identify a priori which assignment method we shall employ. In this case, we'll use DHCP, so there needs to be an IP address available in the DHCP pool with addresses relevant to the subnet on which the VM is to be provisioned.

If an address is available, the process may continue; otherwise, we need to supplement the pool capacity. Supplementation entails either enlarging the size of the current address pool or by allocating an additional subnet from which a pool can be defined. A DHCP server can typically be configured with multiple subnets as "shared subnets" where multiple pools are considered common to a given physical subnet. Once adequate pool sizing has completed, the DHCP server needs to be configured according to the expansion strategy, then it will be ready to distribute an IP address to the pending VM.

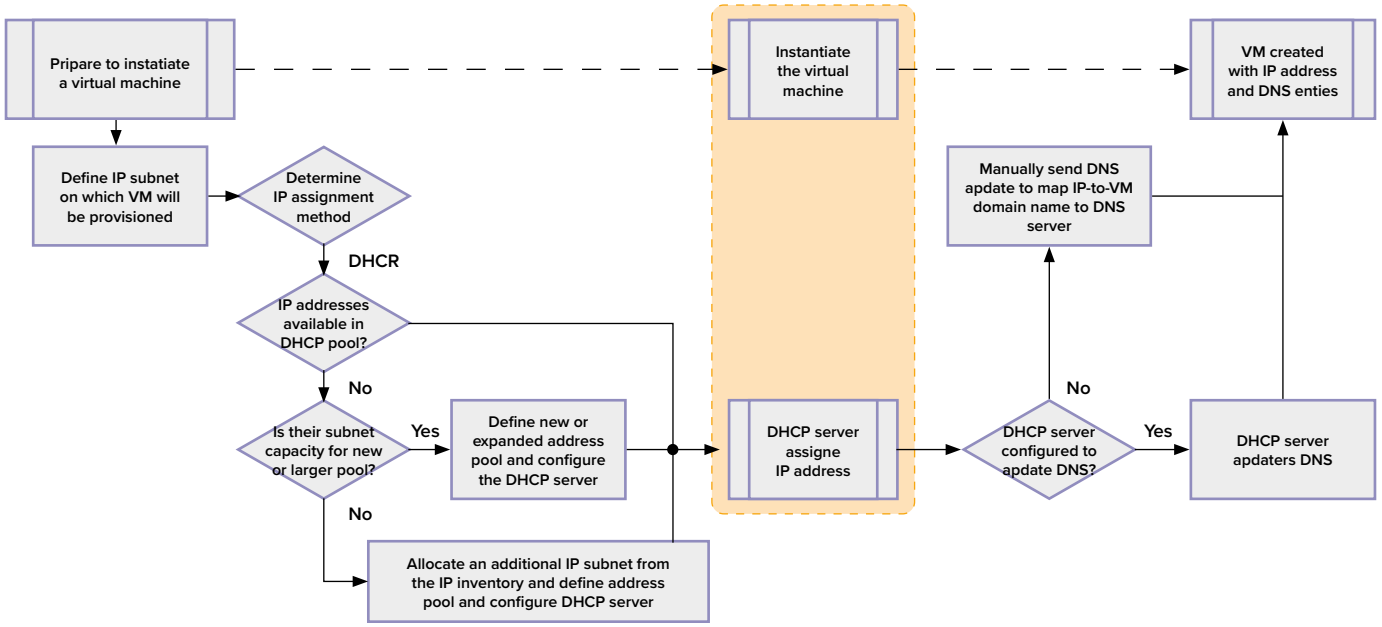

Figure 2: IP initialization using DHCP

The VM can be created and configured to obtain an IP address(es) from DHCP. The DHCP server can be configured to update DNS with the VM's hostname-to-IP address mapping automatically. If the assigned IP address needs to be added to a functional DNS resolution, e.g., service-element2.example.com, this entry will likely require manual entry. First determine the assigned IP address, then create an entry to update DNS to associate the IP address to the function name.

Static Method

If use of DHCP is undesirable, a manual IP address assignment method may be employed. Considering Figure 3, we see the same three macro-level steps at the top of the figure. In this case our assignment method is manual, so we must next identify an IP address for the relevant subnet that can be assigned to the pending VM. If our inventory of subnet IP addresses indicates a free IP address, the chosen IP address should be denoted as used to prevent future erroneous assignment of the address in duplication.

It's a good idea to periodically validate your inventory by comparing it with actual IP assignments, discernable by pinging or otherwise communicating with each host on the subnet. This "discovery" process helps identify those host IP addresses that perhaps were inadvertently not recorded as in-use or those that were unilaterally assigned outside of the process. At minimum, it would make sense to at least ping the IP address to which you plan to assign to the pending VM to verify its vacancy prior to assignment. The assignment of the same IP address to two or more hosts will render neither of them able to communicate.

If an IP address is not available for assignment, allocation of an additional subnet is required. This action too necessitates consulting the IP address space inventory to identify a free subnet that can be assigned. Once assigned and provisioned, an available IP address from the subnet may be recorded then assigned to the VM pending

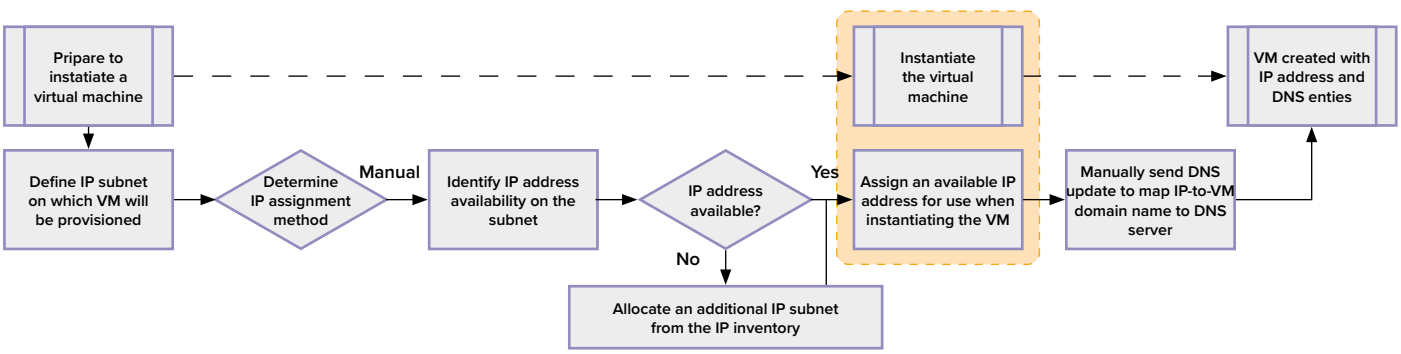

Figure 3: IP initialization using manual assignment

instantiation. This IP address with hostname and/or function name mapping should also be updated in DNS using nsupdate or similar utility.

Both the DHCP and manual IP approach have several potential issues, summarized as follows.

- For DHCP, pools must be monitored proactively so IP addresses are available on demand; configuration of multiple pools or discriminating option parameters can be error-prone and tedious.
- The manual approach requires a per-VM manual step which is error-prone, time consuming and as such not very scalable.
- Both approaches require accurate tracking of IP address space to enable confident expansion of address pools, assignment of IP addresses and allocation of additional subnets.

Public Cloud Method

As mentioned earlier, the assignment of IP addresses within a public cloud infrastructure will generally be controlled by the public cloud platform itself. Some public cloud services hold over an IP address for example, keeping the address unassignable for a time period after its corresponding VM/VNF was destroyed to allow the address to "time out" of any systems in which it was being tracked.

In such a scenario, the IPAM integration process entails tracking the cloud-assigned IP and DNS information upon successful instantiation. Figure 4 shows a simple scenario, where the VM is instantiated and the IPAM system is triggered to poll for the VM's assigned IP address and domain name in order to update the IPAM repository. This cooperative method enables the centralized tracking of IP addresses and DNS names in a holistic database for public clouds where the cloud is authoritative for the address space as well as private clouds and internal networks where the DDI system is generally authoritative.

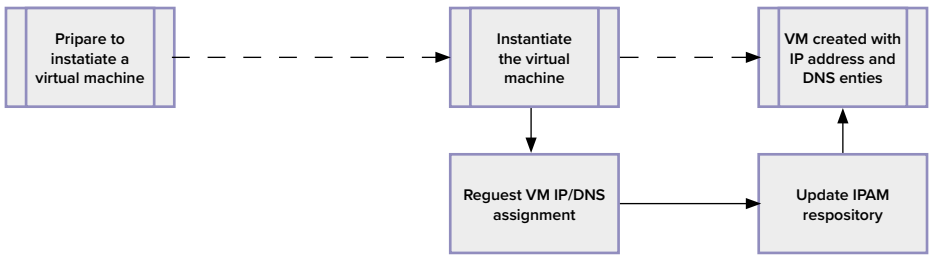

Figure 4: IP initialization within a public cloud

As in the private cloud scenario, discovery of IP address assignments is recommended to align public cloud IP assignments with those tracked within your IPAM system. This is helpful in not only assigning unique IP addresses but for monitoring IP address capacity to enable allocation of additional subnets as needed well before IP address capacity exhausts and VM/VNF creation is rendered impossible.

Public Cloud Method

As mentioned earlier, the assignment of IP addresses within a public cloud infrastructure will generally be controlled by the public cloud platform itself. Some public cloud services hold over an IP address for example, keeping the address unassignable for a time period after its corresponding VM/VNF was destroyed to allow the address to "time out" of any systems in which it was being tracked.

In such a scenario, the IPAM integration process entails tracking the cloud-assigned IP and DNS information upon successful instantiation. Figure 4 shows a simple scenario, where the VM is instantiated and the IPAM system is triggered to poll for the VM's assigned IP address and domain name in order to update the IPAM repository. This cooperative method enables the centralized tracking of IP addresses and DNS names in a holistic database for public clouds where the cloud is authoritative for the address space as well as private clouds and internal networks where the DDI system is generally authoritative.

A Unified Approach

With these various approaches for cloud systems' IP address assignment processes, how can one track and manage IP address space holistically while accommodating all of these different approaches? The use of a robust and flexible IP address management (IPAM) solution such as IPControl™ in conjunction with the Sapphire Cloud Automation Appliance (CAA) from Diamond IP enables the unification of IPAM across these diverse processes.

IPControl is a comprehensive IPAM system that enables you to define your IP space on your terms and provides single-click subnet and IP address assignments via a web graphical interface or via an application programming interface (API). The Sapphire CAA enables the automated creation and deployment of subnets, VM IP addresses and DNS resource records upon instantiation and supports various discovery functions to assure the accuracy of IP address deployments across non-cloud, private cloud, public cloud, and hybrid cloud domains.

Private Cloud Operation

Cygna Labs Diamond IP offers the Sapphire CAA to automate and manage orchestrator functions related to IP address and DNS name assignment. In the context of this white paper, the term orchestrator refers to any northbound system that incorporates DDI functions within its orchestration tasks and includes such systems as ServiceNow, terraform, Ansible, puppet, and other similar platforms. Thus, when a VM is created or destroyed, the orchestrator may be configured to invoke a CAA API call to either assign or free up the corresponding VM IP address and DNS resource records. Figure 5 illustrates this process for a private cloud scenario where IPControl is authoritative for IP address assignment. The orchestrator calls the Sapphire CAA REST application programming interface (API) and the Sapphire CAA shepherds orchestrator requests for IP address and DNS information into the centralized IPControl IPAM database. This provides a fully automated, hands-free mechanism for robust address assignment.

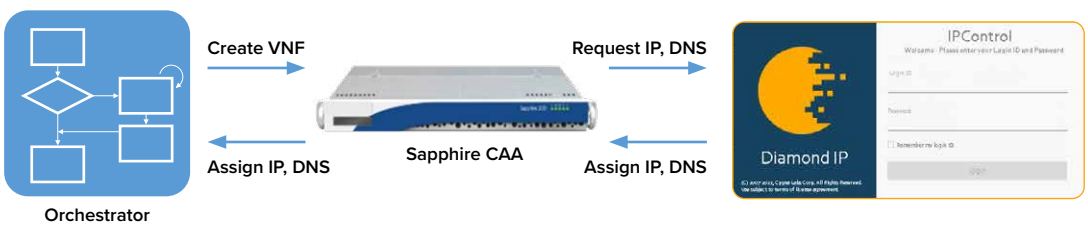

Figure 5: Private cloud IP initialization with IPControl

From an IP address and pool capacity management perspective, IPControl provides periodic monitoring of address utilization. User-definable thresholds enable alerting of administrators of the need to allocate more capacity or even the automated allocation of supplemental capacity without user intervention.

Figure 6 illustrates our basic three step VM instantiation process when used with IPControl. The preparation phase consists of IPControl periodically analyzing DHCP pool capacity, subnet allocations on routers and individual IP address assignments. Impending address pool or subnet IP address capacity exhaustion detection enables hands-free IP capacity management so you can be confident an IP address will be available for assignment when required. Capacity thresholds provide a means to alert administrators with settable severity to provide "info" level alerts at a given threshold level, e.g., 75% or 20 addresses remaining, while providing "warn" and "critical" at more imposing levels accordingly. Threshold triggers can even initiate subnet allocation with a simple API call to the IPControl system.

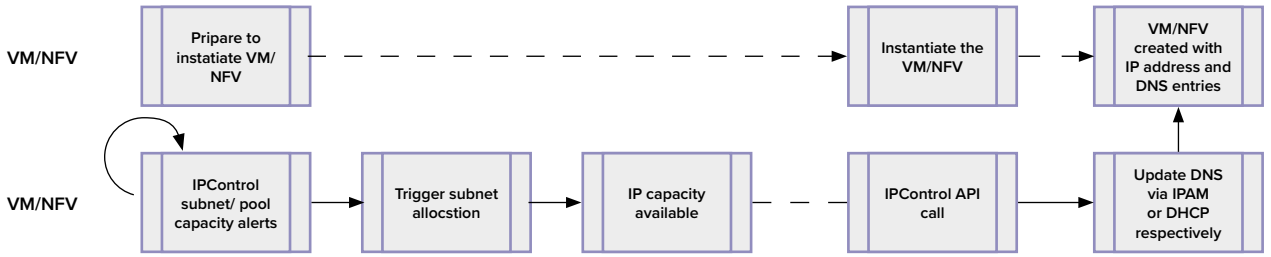

Figure 5: Private cloud IP initialization with IPControl

Subnets allocated via the Sapphire CAA or the IPControl user interface do not require explicit specification of the subnet address. IPControl tracks your overall address space and associated allocations and remaining free space. Automated subnet allocation may use a “best fit” approach and for IPv6 blocks additionally sparse and random allocation methods. In all cases, IPControl tracks allocations and remaining free space so successive allocations can be made easily and accurately with a single mouse click, API call, or via the Sapphire CAA. Subnet templates also enable definition of IP address assignments or reservations. DHCP pools can be templated within subnets to auto-define DHCP pools within subnets for DHCP deployments. Templates not only simplify allocation but promote consistency of address assignments within subnets for administrative and troubleshooting purposes.

IP host discovery via ARP, neighbor discovery or ICMP may also be scheduled on a recurring basis to provide a periodic sanity check on actual network address assignments. Newly discovered IP occupants can be directly imported or saved for review within the planned vs. actual display. This enables synchronization of the IPAM database with the IP network to minimize the possibility of conflicting IP address assignments.

With the preparation phase ongoing in the background to monitor address availability, you can confidently create and destroy VMs dynamically based on your services and capacity needs. Using an orchestrator plug-in, the actual VM creation process obtains an available IP address from IPControl via the Sapphire CAA. The Sapphire CAA automates the process and retains the tracking of IP address assignments within the holistic IPControl repository. IPControl can automate DNS resource record creation as well based on naming policies so that DNS too can be updated without human intervention. The overall process yields a virtually hands-free integration of critical network services initialization within the process of VM instantiation. This saves time, eliminates errors and scales massively.

The IPControl system provides built-in reports for utilization monitoring, administrator and IP address auditing and IP space management. Granular administrator controls enable scoping of visibility and control of the network portion and functions permitted for certain administrators or groups. IPControl is extremely scalable and is in production managing among the largest IP networks on the planet.

Public Cloud Operation

The same IPControl system and Sapphire CAA supports public cloud interaction as well as private cloud to support private-only, public-only or hybrid cloud environments. In the public cloud scenario however, IPControl serves as the IP tracking system based on public cloud IP address assignments.

The Sapphire CAA supports API calls for public cloud operation to allocate or deallocate subnets, instantiate or destroy VMs and to discover subnets and IP assignments. Figure 7 illustrates this basic process flow for instantiating a public cloud VM. Through a single “Create VNF Instance” API call from an orchestrator or automation system, the CAA polls the cloud API for a free IP address, instantiates the VM using an available IP address, then updates IPControl with the IP address and name assignment made by the public cloud. Such “canned” flows provided by the CAA can be modified based on your particular operating requirements. This customization provides superior adaptability of the IP/DNS assignment process.

Discovery functions provide for the collection and comparison of IP addresses from both the cloud system and IPControl. Discovery API parameters define the processing of address assignment differences, options for which include reporting of differences or adding IP addresses to IPControl that are currently assigned only in the cloud and/or deleting IPControl IP addresses that are not assigned in the cloud. In this manner, IP address assignments may be detected and optionally synchronized with those tracked in the IPControl IPAM repository. This enables accurate tracking of IP address assignments within public cloud services, analogously to that provided for non-cloud and private cloud networks.

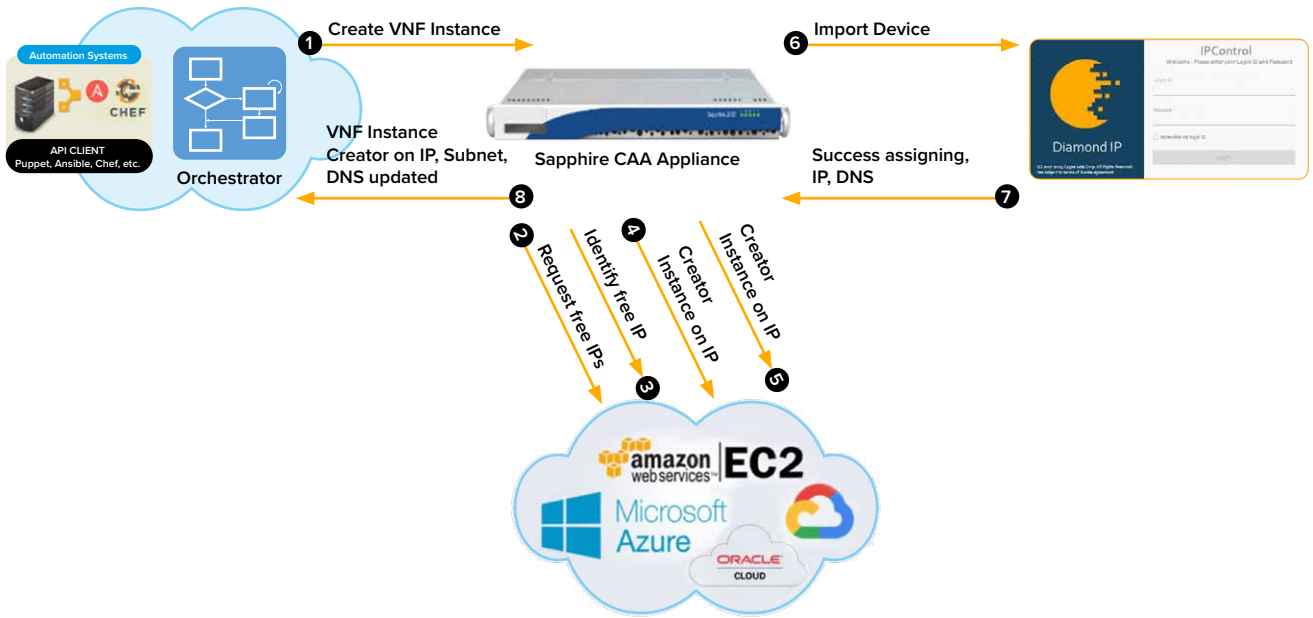

Figure 7: Public cloud IP initialization with IPControl

Beyond Cloud Operation

Automating private and public cloud IP addresses, subnets and DNS information optimizes DDI efficiencies, but many organizations manage network and compute resources beyond those in the cloud. This includes good old branch or remote offices, SD-WAN Internet breakout sites and even good old data centers. The Sapphire CAA provides automation capabilities for these environments as well, enabling you to fully automate DDI across your diverse network estate.

For example, you can create native html pages on the CAA to create a user portal to empower end users to issue requests for IP addresses and DNS information for new devices. For example, a user may request a static IP address for a new printer. A portal page could be created which could trigger a Sapphire CAA flow to assign the next IP address on the subnet where the user is located and provide instant results. This enables IT to provide rapid services provisioning without the need to engage manual IT effort for trivial tasks. The IP information is updated automatically and IT staff can be notified if desired or may review audit logs to spot check address assignments as desired.

IPControl as Orchestrator

Basic DDI functions may require updating of external resources such as external DNS hosting systems. For example, IPControl can manage Akamai DNS as a native DNS vendor like it does ISC, Microsoft and Cisco. Deploying DNS zones and resource records to Akamai DNS services may be accomplished using the Sapphire CAA and our Akamai flows. This enables administrators to manage internal and external DNS holistically within a single system and user interface.

Other DDI events may necessitate external actions as well. For example, the provision of a subnet on a router interface in IPControl can trigger a callout to the Sapphire CAA which in turn may update a router configuration system with the corresponding subnet information, streamlining this provisioning process. IPControl supports several callout events which can be easily configured to trigger CAA flows to update other business or operations support systems, network elements or even send notifications of triggered events.

Core DDI Network Services Benefits

Implementing a disciplined and automated DDI foundation as a core component of your network management and cloud deployment strategy enables you to achieve the following benefits:

- **Faster VM/VNF provisioning** – Whether using DHCP, an orchestrator plug-in to automatically assign a VM IP address, and/or a public cloud platform, the provisioning process need not pause to assign and IP address and update DNS. The Diamond IP solution supports each form of core network services automation for rapid provisioning, supporting the cloud's agility benefit to organizations. And IPControl supports the underlying capacity management functions to provide further assurance of IP address availability during the provisioning process.
- **Improved provisioning accuracy** – Automation via DHCP or the Sapphire CAA provides more accurate provisioning. The requestor for an IP address need only issue the request and leave it to the Sapphire CAA to identify a free address for assignment whether in IPControl or within the public cloud platform. The requestor need not attempt to identify an address on its own and request it; the requestor can request the Sapphire CAA make or obtain the IP address assignment and to maintain the IPAM repository in IPControl.
- **Reduced manual effort** – Automation reduces manual effort, which can reduce time intervals, opportunities for erroneous manual entry and staff costs.
- **Simpler troubleshooting** – With all of your IP address information secured in a robust repository including all assignments and discoveries recorded, the IPControl database serves as the IP address plan of record which provides critical information during troubleshooting. Subnet and IP address templates also promote consistency of IP address assignments which further reduces confusion and aids in rapid troubleshooting.
- **Integrate DHCP/DNS/IPAM (DDI) processes** – As we've seen in discussing the functional integration of assigning IP addresses, updating the IP address repository then adding the IP address-to-name association in DNS, the core network services of IPAM, DHCP and DNS are tightly inter-related. Integrating these services under a single management system facilitates automation and robustness of IPAM data.
- **Centralized DDI for your entire network** – depending on the expanse of your network you likely have non-cloud network components you need to manage, like subnets in remote offices for example. IPControl enables you to manage all of your IP space holistically, integrating the view of cloud IP space and non-cloud IP space through a single pane of glass.
- **Segment administrative authority** – Different administrators may be responsible for certain subsets of the network or subsets of functions like DHCP or DNS for example. IPControl supports very granular administrator roles to define the scope of control accordingly. Even cloud orchestrator systems and the Sapphire CAA can be constrained to the span of their reach when calling the API to just the cloud environment. Likewise, individual administrators or groups can be constrained in multiple dimensions.
- **Virtualized DHCP and DNS network functions** – Cygna Labs Diamond IP offers virtual appliances for DHCP and DNS services, as well as the IPControl centralized IPAM systems itself and the CAA. As your need for elasticity for these core network services dictates, these appliances can be instantiated and destroyed across your VMware, Xen, Hyper-V, KVM, AWS, Azure, or Oracle Cloud infrastructure.
- **Bottom line: lower cost** – Deployment of the Sapphire CAA with IPControl can help you lower costs of cloud administration though automation, less manual staff effort, fewer errors to troubleshoot, and higher customer or constituent satisfaction through rapid, accurate provisioning. This approach mirrors and supports those benefits of the cloud itself with agility, elasticity and lower costs.

Summary

Core DDI network services supply a critical role in the cloud provisioning process. These services automate the assignment of IP addresses to VMs and VNFs and also update DNS for simpler name based location of cloud resources. Without proper management of cloud core network services, you may face an inability to provision VMs when needed due to the lack of available IP addresses just when you need them. In addition, even if an IP address is assigned, improper management may lead to unreachability via the IP address if a duplicate or via name if DNS is not updated correctly. Finally mismanaged IP address space can lead to a capacity reduction and suboptimal address utilization.

Avoid these pitfalls and utilize a disciplined full-cycle DDI system like IPControl with cloud automation features via the Sapphire CAA. The Sapphire CAA with IPControl facilitates secure automation of provisioning, providing and supporting those benefits of the cloud with agility, elasticity and lower costs.

Toll Free: **(844) 442-9462**
International: **+1 (305) 501-2430**
Fax: **+1 (305) 501-2370**

Sales: **sales@cygnalabs.com**
Support: **support@cygnalabs.com**
Billing: **finance@cygnalabs.com**

cygnalabs.com

CYGNA LABS