

# Experimenting with AWS Direct Connect using Chameleon, ExoGENI, and Internet2 Cloud Connect

Paul Ruth and Mert Cevik  
Renaissance Computing Institute  
{pruth,mcevik}@renci.org

**Abstract**—Many scientific research communities and institutions are adopting the cloud as a primary platform to support their computing needs. Rapid adoption of the cloud for scientific computing is a result of the simplicity with which an individual researcher can obtain large amounts of customized compute and storage resources. At the same time, most public cloud providers have rolled out many advanced networking services. Many of these services, like their compute services, are simple to access by any cloud user (e.g. routing between regions and private networking spaces within a cloud). However, it is not possible for most researchers to access expensive low-level, externally facing cloud network services without complicated support by campus IT staff, as well as national and regional network providers. This paper describes how to use Chameleon, ExoGENI, and Internet2’s Cloud Connect service to deploy research experiments that use AWS Direct Connect without requiring a privately owned Direct Connect endpoint or support from local campus IT staff.

## I. INTRODUCTION

Direct low-level network connections between public clouds and scientific facilities can provide the bandwidth needed for large scientific data transfers, low latency required for seamlessly deploying campus facilities across hybrid clouds, and the security of private isolated computing environments. Most public cloud providers have low-level networking services (e.g. AWS Direct Connect [1], Azure ExpressRoute [2], and Google Dedicated Interconnect [3]). However, wide-spread experimentation with these low-level wide-area network connections is limited by access to expensive direct connect endpoints on public clouds as well as complicated campus network configuration requirements. Even if an individual researcher has access to campus staff with the skills required to successfully use these services, it may be difficult to justify the additional staff effort for smaller or short-lived projects. As a result, most research on these services is being performed by a few select scientists or campus IT staff themselves.

Internet2 [4], Chameleon [5], and ExoGENI [6] have recently enabled services that combine to reduce the hurdles that prevent individual researchers from accessing direct connections to public clouds. These services enable a wide array of cloud experiments, as well as provide expanded training opportunities to the next generation of campus IT staff in the use of these otherwise inaccessible services. This paper briefly describes these new services and shows how to combine them in order to deploy your own experiment using AWS Direct Connect.

## II. ENABLING SERVICES

### A. AWS Direct Connect

The AWS networking abstraction is based on the Virtual Private Cloud (VPC). This abstraction allows users to create custom isolated network spaces into which they can deploy virtual machines. VPCs include user specified IP subnets, routers, gateways to the Internet, and can (optionally) be peered to other VPCs or your external data center or other facility.

Peering between cloud resources and external facilities is achieved through the AWS Direct Connect service. This service allows users to add direct L2 network circuits between a virtual router in their VPC and an external router at their campus or datacenter. From the perspective of the campus, a VPC is a peered Autonomous System (AS). Both the virtual router and the campus router use BGP to advertise known routes to each other. Direct connections are similar to *stitched* circuits created by GENI [7] but enable experiments with the wide array of public cloud offerings including scale that is not achievable on existing academic testbeds alone.

### B. Internet2 Cloud Connect

Internet2 has responded to the needs of campus IT and academic researchers who want to use direct cloud connections but cannot justify the cost of their own connections. The Internet2 *Cloud Connect* [8] service comprises a set of shared direct cloud connections at four sites (Ashburn, VA; Chicago, IL; Dallas, TX; and San Jose, CA) available to member institutions. Each site has 10 Gbps shared connections to Amazon, Google, and Microsoft clouds. The connections are managed using the existing Advanced Layer2 Services (AL2S).

Although the Cloud Connect service is a vital piece of our end-to-end solution it does not solve the problem on its own. Researchers would still need to provision circuits across their regional provider and from their campus edge to the target facility and configure a BGP router.

### C. ExoGENI

ExoGENI is one of two testbeds deployed as part of the GENI initiative. The strength of ExoGENI is its ability to provision (i.e. *stitch*) L2 circuits across I2 AL2S as well as many regional network providers. As part of this work, ExoGENI has been augmented to provision stitched L2 circuits between Internet2 Cloud Connect endpoints and any existing ExoGENI resource.

ExoGENI *stitchports* are reusable pre-provisioned circuits that cross regional networks and campuses ending at existing

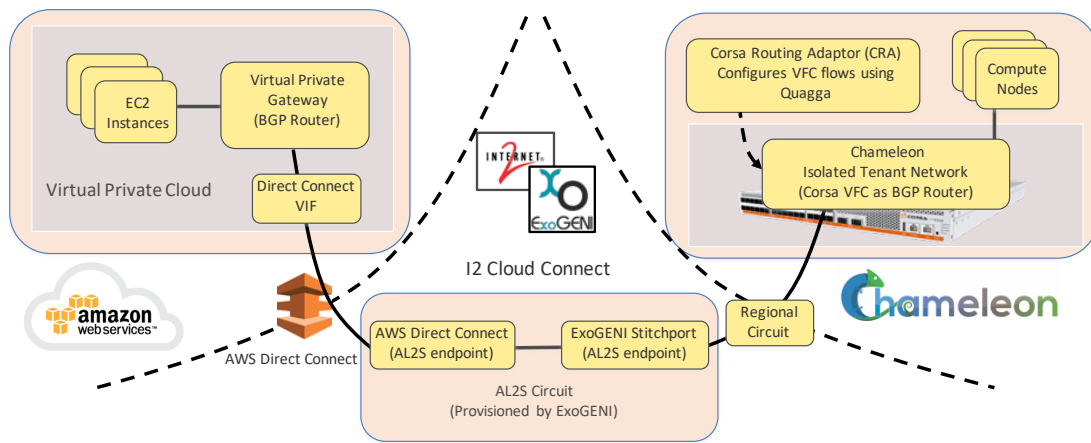


Figure 1. End-to-end experiment deployed across Chameleon, ExoGENI, and Amazon using Internet2's Cloud Connect service.

external facilities including other testbeds. Dozens of *stitchports* have been provisioned to various facilities since the inception of ExoGENI. Any existing ExoGENI stitchport can easily be connected to any Internet2 Cloud Connect endpoint. For example, this work utilizes a set of ExoGENI stitchports that connect directly to the Chameleon testbed. Chameleon users can deploy baremetal cloud resources connected to an ExoGENI stitchport then use ExoGENI to connect that stitchport to any Internet2 Cloud Connect endpoint.

#### D. Chameleon

Chameleon is a large-scale, deeply reconfigurable experimental testbed built to support Computer Science systems and networking research. Chameleon supports a bare metal reconfiguration of compute systems giving users full control of the software stack including root privileges, kernel customization, and console access. In addition, Chameleon provides isolated user programmable network hardware using OpenFlow and Corsa DP2000 series switches [9]. Corsa switches are natively sliceable into multiple isolated Virtual Forwarding Contexts (VFCs). Chameleon users can program networks using their own OpenFlow controllers connected to an isolated VFC. Optionally, ExoGENI stitchports can be connected directly to the VFCs enabling user experiments spanning Chameleon and any resources accessible via ExoGENI.

Many Chameleon users experiment with networks that they control using standard OpenFlow controllers. However, users can also deploy a Corsa Routing Adapter (CRA) which is a powerful tool that simplifies programming a VFC to be a router. The CRA is deployed on an additional Chameleon node and acts as a connector that translates rules between a software Quagga router and the VFC. The result is a VFC that is programmed using the standard Quagga interface. The Quagga/CRA controlled VFC can even be configured with BGP and peered to an external router.

### III. END-TO-END AWS DIRECT CONNECT EXPERIMENTS

Figure 1, shows the complete end-to-end experiment spanning Chameleon, ExoGENI, and Amazon (via Internet2 Cloud

Connect). Users can deploy EC2 instances in their own VPC and connect nodes in an isolated Chameleon tenant network. Both the VPC and the Chameleon network are configured with BGP router that advertise routes to each other. The VPC includes a Virtual Private Gateway configured as a BGP router while the Chameleon isolated network is a VFC programmed by Quagga/CRA to be a BGP router. The routers are connected by L2 circuits provisioned by ExoGENI across AL2S and the Chameleon regional providers.

### IV. CONCLUSIONS

Experimenting with AWS Direct connect is challenging for a typical researcher. New services from Chameleon, ExoGENI, and Internet2 combine to enable AWS Direct connect experiments.

### ACKNOWLEDGMENT

Results presented in this paper were obtained using the Chameleon and ExoGENI testbeds supported by the National Science Foundation. This material is based upon work supported by the NSF awards 1743313 and 1836715.

### REFERENCES

- [1] Amazon, "AWS Direct Connect," <https://aws.amazon.com/directconnect/>.
- [2] Microsoft, "Azure ExpressRoute," <https://docs.microsoft.com/en-us/azure/expressroute/expressroute-introduction>.
- [3] Google, "Google Dedicated Interconnect," <https://cloud.google.com/interconnect/docs/concepts/dedicated-overview>.
- [4] Internet2, "Internet2," <https://www.internet2.edu/>.
- [5] K. Keahey, P. Riteau, D. Stanzione, T. Cockerill, J. Mambretti, P. Rad, and P. Ruth, *Chameleon: A Scalable Production Testbed for Computer Science Research*, 05 2019, pp. 123–148.
- [6] I. Baldin, J. S. Chase, Y. Xin, A. Mandal, P. Ruth, C. Castillo, V. Orlikowski, C. Heermann, and J. Mills, "Exogeni: A multi-domain infrastructure-as-a-service testbed," in *The GENI Book*, R. McGeer, M. Berman, C. Elliott, and R. Ricci, Eds. Springer, 2016, pp. 279–315. [Online]. Available: <http://dblp.uni-trier.de/db/books/collections/MBER2016.html#BaldinCXMRCOHM16>
- [7] R. McGeer, M. Berman, C. Elliott, and R. Ricci, Eds., *The GENI Book*. Cham: Springer International Publishing, 2016. [Online]. Available: <http://dx.doi.org/10.1007/978-3-319-33769-2>
- [8] George Loftus, "Internet2 Announces Production Cloud Connect Service General Availability," <https://www.internet2.edu/blogs/detail/17059>.
- [9] Corsa, "Corsa DP2000 Product Family," <https://www.corsa.com/products/>, 2018.