CloudJoin

Experimenting at scale with Hybrid Cloud Computing

Jack Brassil | ***Irene Kopaliani***

Dept. of Computer Science | Office of Information Technology

PRINCETON UNIVERSITY

Cloud Research Infrastructures

- US academic computing & networking systems research community relies on computing testbeds including *CloudLab*, *Chameleon Cloud*, *Jetstream*, *OpenCloud*
- CRIs designed to
  - support experiment isolation, reproducibility, information sharing, investigation of future clouds
  - Provide research community meeting place
  - extend campus laboratory and Research Computing systems
- CRI developers learn by building and operating
- Experimenters report high degree of satisfaction with CRIs

PRINCETON UNIVERSITY

CRI challenges

- Expensive to sustain and refresh equipment investment
- Limited resources to invest in ease-of-use, monitoring/debugging tools, experimenter training
- Difficult to federate with other infrastructures
- Suffer busy period congestion

PRINCETON UNIVERSITY

Commercial compute clouds

- Offer massive general purpose and specialty computing and storage resources
- Increasingly focused on providing
  - software and services (transcription, translation, video analysis)
  - Variety of software platforms/abstractions (containers, serverless)
  - Hosting large scale science datasets
  - Vertical application support (healthcare, finance)
- Will continue to expand hardware and software services

Commercial compute clouds

- Multiple barriers have slowed academic *research* use adoption
  - Relative costs, cost transparency
  - Moderately steep learning curve
  - Relative unattractiveness for grant support
  - No direct control of provider infrastructure
  - Availability and familiarity with on-campus laboratory, department and institution research computing facilities

CloudJoin

- Recognizes complementary properties of CRI, on-campus computing, & commercial cloud systems
  - use your private resources first, tap commercial cloud as needed
- Explores research & education uses of hybrid cloud architectures
- Focus on use cases not well served by on-campus or CRI alone
  - Large-scale systems experiments spanning local and cloud resources
  - Specialty cloud hardware or services
  - Short term experiments/class projects

PRINCETON UNIVERSITY

Research challenges

- Successful hybrid cloud experimentation places demands on infrastructure development *and* experiment design
  - Connecting local resources to commercial clouds
    - Experiment-level (no infrastructure change required) vs. infrastructure-level
  - Placement of experimental resources
  - Tools for monitoring large-scale experiment operation and behavior
  - New approaches to maintain
    - Performance
    - Isolation
    - Reproducibility
    - Collaboration and data sharing

Key contributions

- Demonstrate hybrid Google Cloud Platform (GCP) and CloudLab infrastructure
- Explore limits of experiment-level connectivity
- Demonstrate utility of commercial quality GCP experiment monitoring for usability, scale, tool integrations
- Develop and share experiment tools, best practices

PRINCETON UNIVERSITY

Experiment-level connectivity

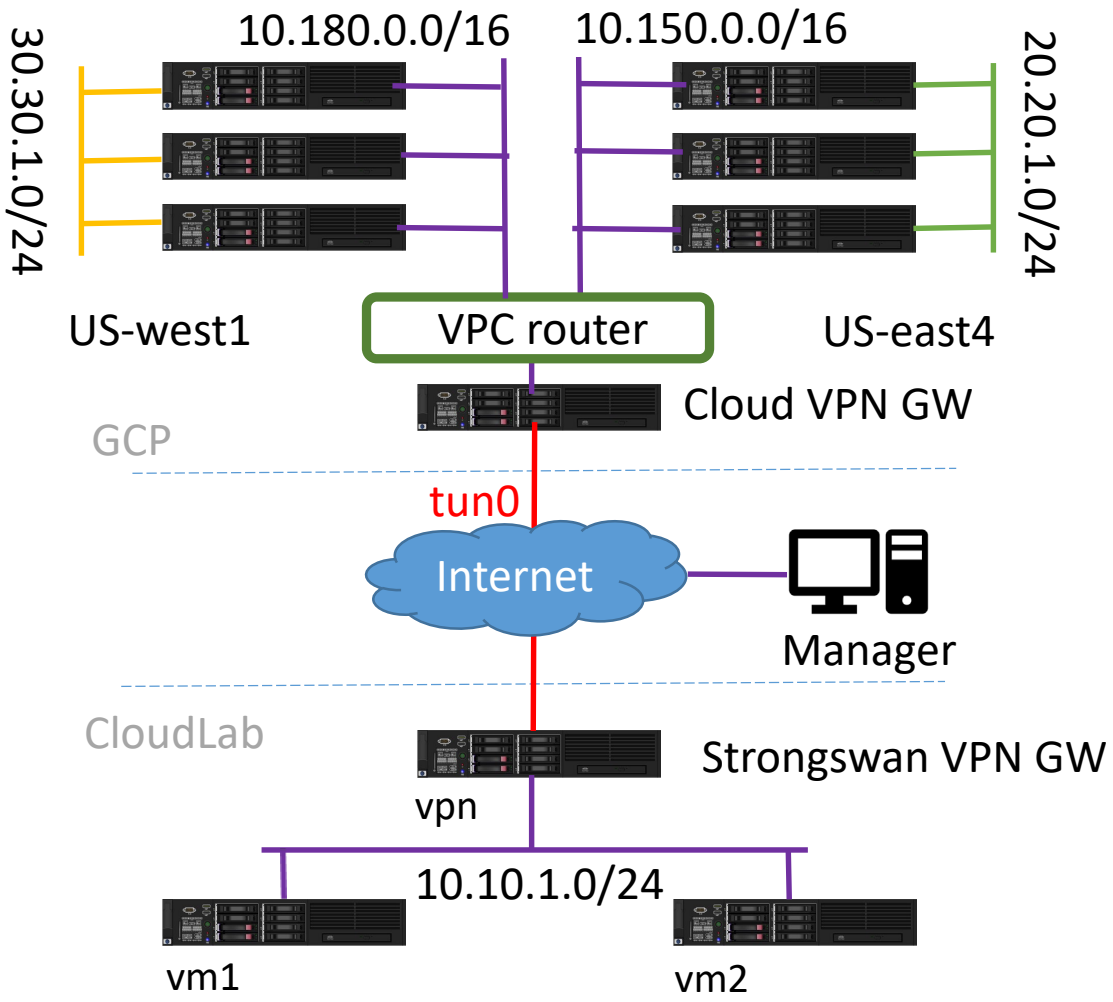

- Desktop
  - Preliminaries (accounts, software, etc)
  - Experiment Manager
- GCP
  - VPC spanning regions
    - VPC router connects subnets
    - Cloud VPN for CloudLab connectivity over public internet
- CloudLab (Clemson)
  - Vanilla experiment topology
  - Dedicated bare metal node or VM for Strongswan VPN GW

Experiment design considerations

- Performance
  - Software VPN throughput limitations
    - Bare metal (2.6 GHz Xeon Gold 6142); CloudLab ⇔ GCP: 750 Mbs
    - VM; GCP ⇔ Cloudlab 476 Mbs (upstream), 250 Mbs (downstream)
  - Latency
    - Roughly equal to public internet across tunnel
    - Pick GCP site locations to minimize latency (!= geography) and cost!
- Observations affecting isolation, reproducibility
  - Multiple tunnels possible to increase performance
  - HA tunnels an option for overcoming tunnel disruptions
  - Control plane communications over public internet
  - Reduced isolation

PRINCETON UNIVERSITY

Establishing data plane connectivity

- Instantiate cloud-side experimental resources (Cloud SDK)
  - Create Virtual Private Cloud (VPC), cloud-side instances, VPN gateway
  - Retain cloud-side connectivity parameters: destination subnets, VPN gateway public address, keys, etc
- Instantiate CloudLab-side experimental resources (portal, geni-lib)
  - Create dedicated VPN node, configure VPN with cloud-side parameters
  - Configure manual routing on nodes for VPC access
  - Retain cloudlab-side connectivity parameters
- Create encrypted VPN tunnel (illustrative script)
  - Select and start cloud VPN gateway: e.g., gcloud compute vpn-tunnels create [params]
  - Setup VPC routing via tunnel : e.g., gcloud compute routes create
  - Start CloudLab VPN gateway service

Monitoring large-scale CloudJoin experiments

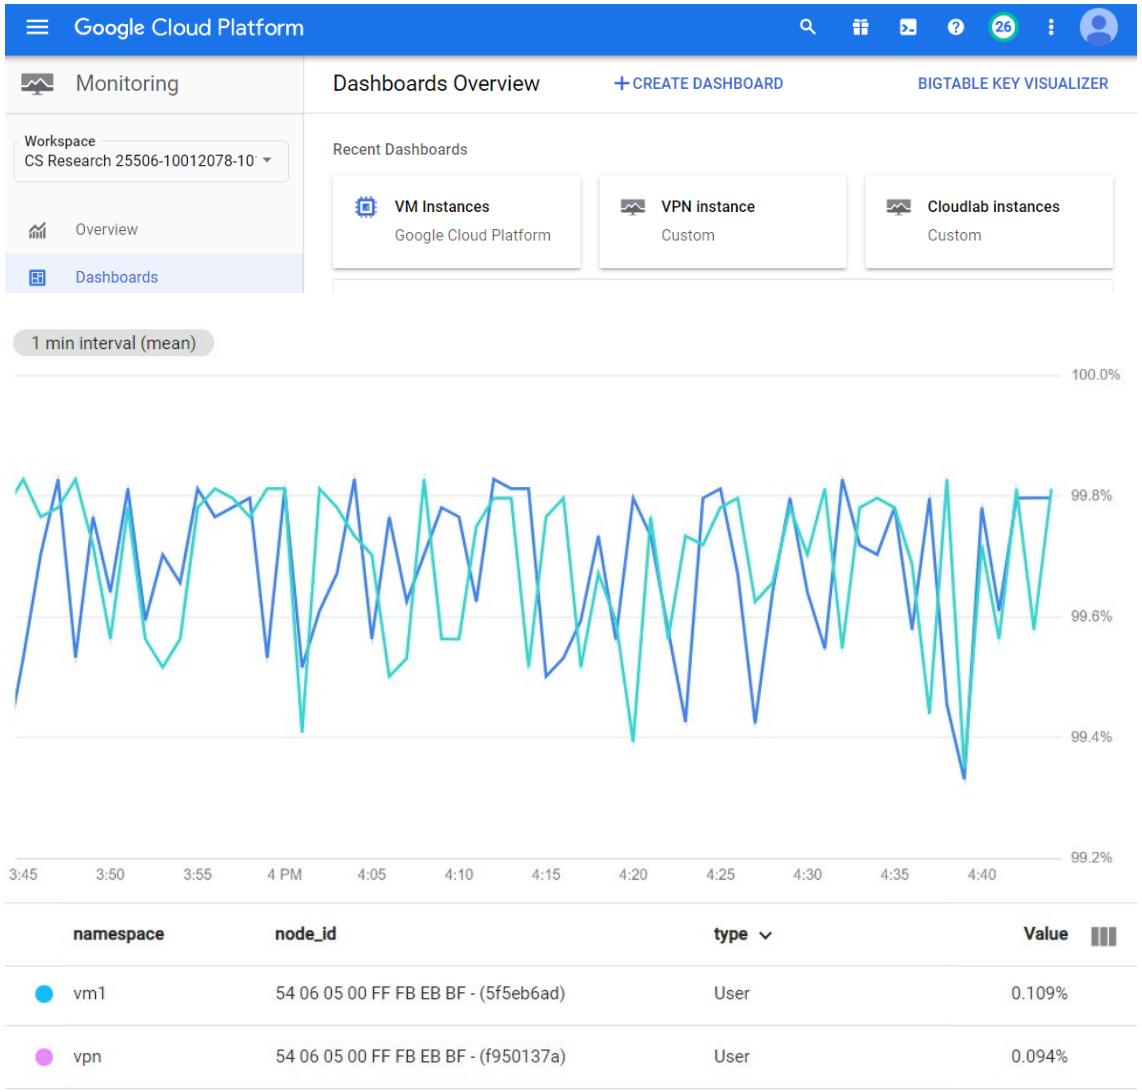

- Desktop
  - Visualize experiment behavior
- GCP
  - *Stackdriver* monitoring API
  - Run stack-driver service agent on instances (*collectd*)
- CloudLab (Clemson)
  - Run BindPlane (*google-fluentd*) monitoring agents on all nodes

Getting started

- Preliminaries – establish accounts, software, familiarity, etc
- CloudLab tools that will help
  - profiles (cloudjoin*N*)

    Description: instantiate simple topology
  - images (cloudjoin*N*.vm1, cloudjoin*N*.vpn)

    Description: CloudJoin-GCP VPN on Ubuntu 16.04/18.04 with strongswan, google-cloud-sdk, BindPlane monitoring, assorted tools
- CloudJoin resources
  - Learn more at https://www.cs.Princeton.edu/~jbrassil/public/projects/cloudjoin

PRINCETON UNIVERSITY

Conclusion

- Hybrid cloud is a viable approach to sustain and enhance investments in CRI and on-campus resources
- Impact: Your experiment artifacts can be more visible and transferable to a broader audience of cloud users
- Next step – infrastructure-level connectivity
  - Characterize large-scale experiment behavior over Direct Cloud Connect
  - Develop tools and best practices for spanning experiments from your on-campus computing resources

Acknowledgement: This material is based upon work supported by the NSF under Grant No. CNS-1923692

PRINCETON UNIVERSITY