

APRICOT 2026 Fellowship Report

Full name	HASAN-UZ-ZAMAN ASHIK
ID Number	20427
Organization	Pico Public Cloud Ltd.
Country	Bangladesh
Fellowship program	Practical Virtualization with Hybrid Strategies APRICOT Fellows Soft Skills Workshop
Location	Jakarta, Indonesia.
Dates	05 Feb 2026 – 08 Feb 2026

Executive Summary

I attended the Practical Virtualization with Hybrid Strategies masterclass delivered by the NSRC team, along with the APRICOT Fellows Soft Skills Workshop. The technical sessions helped me think more clearly about how to choose virtualization approaches that work well for both lab testing and production use. By completing all labs using the shared AWS environment, I strengthened my understanding of secure cloud operations, cost visibility, and practical ways to control spending. I also gained hands-on familiarity with Proxmox (KVM) and the main storage models we deal with in real deployments—block, file, and object storage.

In the Soft Skills Workshop, I practiced organizing my message in a clear structure: what the problem is, why it matters, what solution to propose, and why it is the best option. I also learned how decisions and standards are shaped through community and policy processes. Overall, the workshop improved my confidence in speaking English more naturally and presenting technical ideas in a simple, professional way. I plan to apply these technical and communication skills within my team at Pico Public Cloud Ltd., and also share them through local NOG communities to improve deployment efficiency, cost awareness, and internal knowledge sharing.

Background and Objectives

I work as a DevOps Engineer at Pico Public Cloud Ltd. My main responsibilities include building and maintaining internal network monitoring and reporting dashboards, and integrating SaaS products into our cloud environment. Our platform is based on OpenStack running on Kubernetes, with Proxmox used for internal workload virtualization needs. For distributed and resilient storage, we rely on Ceph.

In day-to-day operations of running a public cloud, we face practical challenges such as IOPS degradation on storage nodes, using cloud-init properly to match different customer requirements, and supporting on-premises to cloud migration work. These issues were the main reasons I joined the masterclass—so I could improve my understanding of the technologies underneath and bring back better operational practices for automation, storage, and deployment.

Objectives

- To gain a deeper understanding of how Ceph works internally and how to operate it more efficiently.
- To learn how experienced professionals design and manage virtualization environments in production.
- To understand practical approaches to automation in cloud and hybrid environments.
- To improve my ability to select the right tools and strategies for optimizing our cloud infrastructure.

Masterclass Overview

Cloud Foundations and Security

During the initial sessions, I refreshed and clarified my understanding of core cloud service models such as SaaS and IaaS, and how responsibility shifts between provider and customer. We revisited the building blocks of cloud infrastructure — compute, storage, and networking — and discussed when to use block, file, or object storage depending on workload needs.

The networking discussions were particularly useful, especially around VPC design, Elastic IP usage, load balancing, DNS, and how outbound traffic costs can quietly increase overall cloud expenses. I also strengthened my understanding of practical cloud security measures, including access control, IAM hygiene, and why SSH key-based authentication is essential instead of password-based access.

Automation and Deployment

One of the most practical parts of the workshop was working with Terraform to provision infrastructure in AWS. This helped reinforce the importance of Infrastructure as Code for consistency and repeatability.

We also explored Cloud-Init in more depth and saw how it can automate instance configuration during the first boot. This was valuable for me because it directly relates to reducing manual configuration steps and improving deployment speed in production environments.

Virtualization Operations

The sessions on virtualization helped connect theory with real operational practices. We reviewed how hypervisors abstract hardware resources and how Proxmox with KVM manages virtual machines in a cluster environment.

Practicing live and cold migration gave a clearer understanding of workload mobility and high availability. We also discussed the practical differences between virtual machines and containers in Proxmox, including why containers may not always be suitable for multi-tenant customer workloads due to shared host dependencies.

Storage

We compared different storage approaches, including block, file, and distributed storage systems, and when each is appropriate. Hands-on work with LVM helped reinforce how flexible storage allocation works at the system level. Learning about Linstor introduced another approach to replicated block storage across nodes.

The Ceph session was particularly important for me. Understanding its architecture and observing a failure simulation gave practical insight into how distributed storage maintains resilience and why proper cluster design and monitoring are critical.

Observability and Governance

The monitoring sessions focused on gaining real visibility into infrastructure health. Using Prometheus and Grafana to visualize metrics showed how proactive monitoring improves troubleshooting and operational awareness.

We also discussed cloud cost management and how easily unnecessary resources or poor design decisions can increase expenses. The data risk management session highlighted the importance of backup, replication, and protecting critical workloads.

Finally, we learned how to evaluate when to use on-premises infrastructure, public cloud, or a hybrid approach based on cost, scalability, compliance, and operational control. This strategic perspective was valuable beyond just technical implementation.

Day-wise Activities and Labs

Day 1 (Thursday)

The first day focused on strengthening our foundation in cloud computing concepts and service models. We discussed how cloud platforms provide core components such as compute, storage, networking, and managed services, and how responsibility changes between SaaS and IaaS models. Storage types like block, file, and object were compared, including how distributed systems like Ceph support these models. There was strong emphasis on architecture design, security, availability, and cost awareness, along with the importance of Infrastructure as Code using Terraform.

In the AWS lab sessions, we practiced creating secure IAM users with MFA, launching EC2 instances, configuring security groups, and working with S3 storage. The cloud security lab helped us identify common risks such as public S3 exposure and overly permissive IAM policies. Through the Terraform lab, we provisioned infrastructure using code and understood how repeatable deployments reduce human error and improve consistency.

Day 2 (Friday)

Day 2 shifted toward virtualization, automation, clustering, and storage operations. We reviewed how hypervisors abstract physical resources to create isolated virtual machines. In the Proxmox lab, we created cloud-image templates, deployed VMs using Cloud-Init, and validated connectivity, which showed how automation reduces manual configuration work.

We explored Cloud-Init in both AWS and Proxmox environments and understood how first-boot automation simplifies large-scale deployments. The clustering session clarified quorum concepts and high availability design. During the cluster lab, we joined nodes, verified cluster health, and performed live VM migration, which demonstrated workload mobility with minimal interruption.

The storage session covered both local and shared storage models. Practicing LVM helped reinforce how flexible storage allocation works in real environments. We were also introduced to Linstor and saw how replicated block storage across nodes supports resilience and VM mobility in clustered environments.

Day 3 (Saturday)

The final day focused on distributed storage, cost visibility, data protection, monitoring, and infrastructure strategy decisions.

By observing simulated failures and automatic recovery, I gained a clearer understanding of how distributed storage maintains resilience and why monitoring cluster health is critical in production environments.

I was able to identify one AWS bucket role permission related issue and reported to instructor earlier so that they can fix it. This helped others not to get the same issue when doing the lab.

In the cloud cost management session, we analyzed usage patterns and spending trends, which highlighted the importance of visibility, tagging, and proactive cost control. The data risk management discussion reinforced the need for versioning, replication, and least-privilege access policies when handling critical workloads.

The monitoring lab introduced practical observability using Prometheus and Grafana. By integrating monitoring with Proxmox and Ceph, we gained better visibility into infrastructure performance and health.

We also reviewed the differences between containers and virtual machines in Proxmox, understanding where containers are suitable and where stronger isolation through VMs is preferable.

The workshop concluded with a discussion on choosing the right infrastructure strategy — whether on-premises, public cloud, or hybrid — based on cost, compliance, scalability, and operational complexity. This strategic perspective helped connect the technical sessions to real business decisions.

Fellows Skills Workshop (Soft Skills)

The Soft Skills workshop presentation called *The Art of Connection* changed the way I look at public speaking. I always thought speaking well was mainly about having strong content, but I realized it is equally about how you deliver the message and how you connect with people.

One thing that really stayed with me was the discussion about nervousness. The trainer compared adrenaline to network congestion — if it builds up without control, everything slows down or crashes. That example made a lot of sense to me as a technical person. I understood that nervousness is normal,

and instead of fighting it, I should learn how to manage it. Simple things like breathing properly, preparing clearly, and knowing the flow of my talk can make a big difference.

Another important lesson was about mindset. Confidence does not come automatically. It grows when we shift our focus from “How do I look?” to “Is my message useful for the audience?” When I think about helping people understand something, I feel less pressure. That small change in thinking can reduce fear a lot.

We also practiced how to speak more clearly. I learned that small changes — like pausing at the right moment, emphasizing important words, and not speaking too fast — can make technical explanations easier to follow. As a non-native English speaker, I also practiced speaking more smoothly instead of word by word. That gave me more confidence.

During the group activity, our team had to present a policy topic chosen randomly. We had to explain what the problem is, why it matters, what solution we propose, and why our solution is the best one. This simple structure helped me see how important logical thinking is when presenting ideas. I also learned that teamwork matters a lot. When everyone shares their thoughts and feels included, the final result becomes much stronger.

Overall, this workshop did not just teach theory. It gave practical tools that I can use in meetings, presentations, and technical discussions in my daily work.

Action Plan After Returning

After returning from APRICOT, my first step will be documenting the key things I learned and the important lab practices in our internal knowledge base. This will help our team have a clear reference for topics like virtualization, Cloud-Init usage, Ceph operations, Terraform workflows, and monitoring setup.

I also plan to start developing reusable Terraform modules for our OpenStack infrastructure, such as VM provisioning, networking, and security groups. The goal is to gradually reduce dependency on manual GUI-based deployment and move toward more consistent, code-driven infrastructure management.

Another focus area will be creating standard Cloud-Init templates for common use cases like web servers, application servers, and monitoring nodes. This should help us speed up VM provisioning and reduce repetitive configuration work after deployment. This depends on customer to customer use case.

On the storage side, I plan to review our current Ceph cluster performance more closely and improve monitoring visibility using Prometheus and Grafana dashboards. Better insight into IOPS, OSD health, and replication status will help us troubleshoot faster and avoid performance surprises.

I also want to explore ways to provide better cost visibility for our public cloud customers, similar to how AWS Cost Explorer works, so customers can understand their usage and avoid unexpected charges.

Within the team, I will arrange a small internal knowledge-sharing session to share key lessons from the workshop. This will help us align on best practices and improve collaboration between DevOps and Operations.

Outside the company, I plan to write a few blog posts about what I learned. Sharing these experiences may help others who are working with similar technologies. I also plan to arrange similar soft skill development program within my team so that we master the tips and tricks shared in the program.

Acknowledgements

I would like to express my sincere gratitude to the APRICOT 2026 organizers, NSRC instructors, and the Fellows Workshop facilitators for arranging such a well-structured and practical learning experience. The masterclass and soft skills sessions were highly valuable, combining hands-on technical labs with real-world operational insights.

I am especially thankful to the trainers for their clear explanations, patience during labs, and willingness to share real production experiences. The opportunity to collaborate with participants from different countries also broadened my perspective. This experience will directly contribute to improving our practices at Pico Public Cloud Limited.