EUROPEAN COMMISSION

Directorate-General for Communications Networks, Content and Technology

CNECT.E – Future Networks
**CNECT.E – Future Networks**

Epics for the Fediversity April MVP - Tweag, a Modus Create Company

Epics

1. Implement the collaboration between the Nixos VM and the Identity Management architectural components.
2. Implement the collaboration between the Orchestrator and the Nixos VM architectural components.
3. Implement the collaboration between the Orchestrator architectural component and the Virtualization architectural group of components.
4. Implement the collaboration between the Nixos VM and the Storage architectural components.
5. Implement the collaboration between the Virtualization and the S3 storage architectural components.
6. Implement the collaboration between the Orchestrator and the S3 storage architectural components.
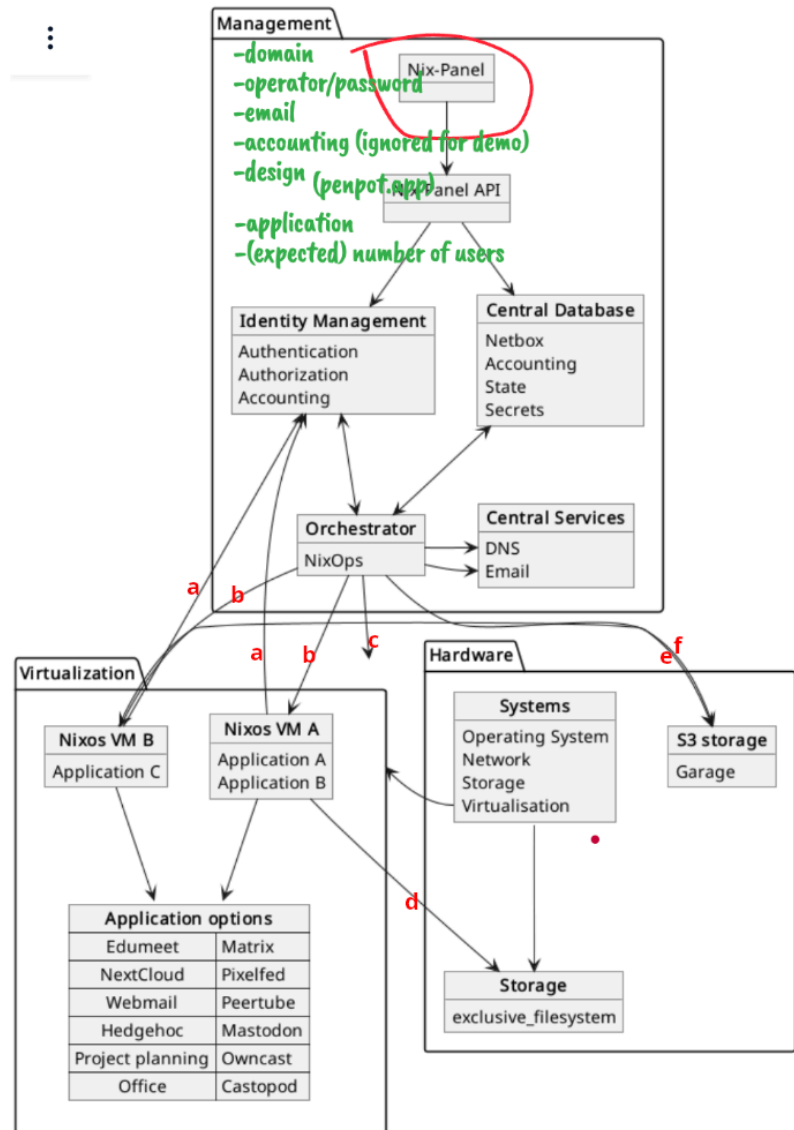
Epics details

1. Implement the collaboration between the Nixos VM and the Identity Management architectural components.

   I don't know what this arrow means. It means we can create an admin account for the services and that my configurations know how to get this information from the hosting provider's databases.

   It is not within the scope of the April MVP because the Identity Management architectural component is not included.

2. Implement the collaboration between the Orchestrator and the Nixos VM architectural components.

   We can deploy NixOS configurations onto existing VMs.

   This comprises several layers:

   1. The deployment software must know how to do this. This is the perimeter of NixOps4 (Robert), which already does the job well in this situation.
   2. We know how to inform NixOps4 about what it should be doing and what basic configurations to write for the provided VMs. In the past months (particularly with the demo in November), we have shown that we know how to do both.
   3. We have well-working multi-machine configurations of Fediverse services. However, this point needs more work. We are only at the proof-of-concept level, and the multi-machine handling is still relatively poor (we are currently working on it).

3. Implement the collaboration between the Orchestrator architectural component and the Virtualization architectural group of components.

   This epic is not Tweag, a Modus Create company's responsibility, but of OID.

   We can provision VMs based on a configuration, but this is just not possible at the moment. NixOps4 does not have a Proxmox provider (Proxmox is the software provided by Procolix to manage VMs; a provider is NixOps4's way of interacting with the real world).

Our project coordinator mentioned in our meeting, "Fediversity refinement session deliverable 2," from Thursday, February 13, that we would not expect anything of this arrow by April and would just use pre-deployed VMs.

We have a workaround solution that can be deployed to cover the demo in April 2025.

4. Implement the collaboration between the Nixos VM and the Storage architectural components.

This epic is not the responsibility of Tweag, a Modus Create company, but of OID in collaboration with NixOps.

It just means that whoever handles the hardware (Procolix, in our case) needs to provide shared storage for the VMs. ProcoliX already does that, and we think it's called "Linstor storage."

We don't think there is much to do on our end.

5. Implement the collaboration between the Virtualization and the S3 storage architectural components.

The Fediverse services rely on external S3 storage for large objects (e.g., pictures, videos, etc.)

This is the case for the three services we have worked on. It can sometimes be clunky, mainly because it currently relies on a Garage instance on the same machine as the service.

We'll separate the service machines from the S3 machines and will have that done before April.

6. Implement the collaboration between the Orchestrator and the S3 storage architectural components.

It is unclear to us what this is. It could be two things:
1. The orchestrator knows to inform the external S3 storage of new bucket creation
2. We can deploy resilient S3 storage from the orchestrator, and this is between the Orchestrator and the OID partner team members.

In Case 1.: Do we ask the hosting provider (e.g., Procolix) to provide one big S3 storage (e.g., Garage) that all the operators will share, or do we only ask the hosting provider to provide disk space, and we spin up an S3 cluster as part of all the VMs that we deploy?

We should rely on the second option for April's MVP since we can already deploy Garage as part of our experiments.

This is closer to the second option specified above, and there are **two main limitations**:
- It is not **resilient and robust**, but we are discussing a prototype, not a final product.
- For now, it **is deployed on the same machine as the servic**e, which is quite a limitation. **This is a work in progress.**