NCShare offers both web-based (Container Manager, Open On Demand) and Linux console-based interfaces for compute services to researchers and students across more than 20 participating institutions in North Carolina. Since its inception, the project has sought to take advantage of federated identity, both to minimize friction for researchers and students and to minimize management overhead for NCShare staff. Because applying federated SSO to console interactions, while possible (cf. Project Moonshot), introduces significant management overhead *and* significant end-user complexity, the project opted for a blended IAM strategy that combined lightweight multi- and bi-lateral SAML federation for access to Container Manager with an integration with GitLab.com's OIDC IDP for access to Open On Demand and a custom provisioning integration between gitlab.com and an NCShare-maintained OpenLDAP directory service to provide both authentication (through directory-distributed SSH public keys provisioned via an API exposed by gitlab.com) and POSIX account services to the Linux hosts comprised by the NCShare slurm cluster. This initial IAM model worked well in the initial phases of the project, but suffered from some shortcomings that became increasingly apparent over time, including:

- risk associated with the dependency on “free” services from GitLab.com, which could be monetized or decommissioned without notice by the vendor
- complexity of user onboarding, as researchers had to establish and then manually link their GitLab.com identities with their institutional identities through a custom interface provided by NCShare, and use GitLab.com's IAM service to manage SSH keys for access to NCshare Linux hosts
- confusing IAM fragmentation, as end users would use home site credentials to access Container Manager and GitLab.com credentials to access Open On Demand
- scaling difficulties associated with onboarding new web-based interfaces inside NCShare, since each web-based interface in NCShare would require explicit integration with every non-InCommon-participant school's SSO service, imposing undue overhead on both institutional and NCShare IT staff for any new deployments
- severely limited authorization flexibility due to the lack of a clear mechanism for managing group memberships or role assignments for distributed end users of NCShare services

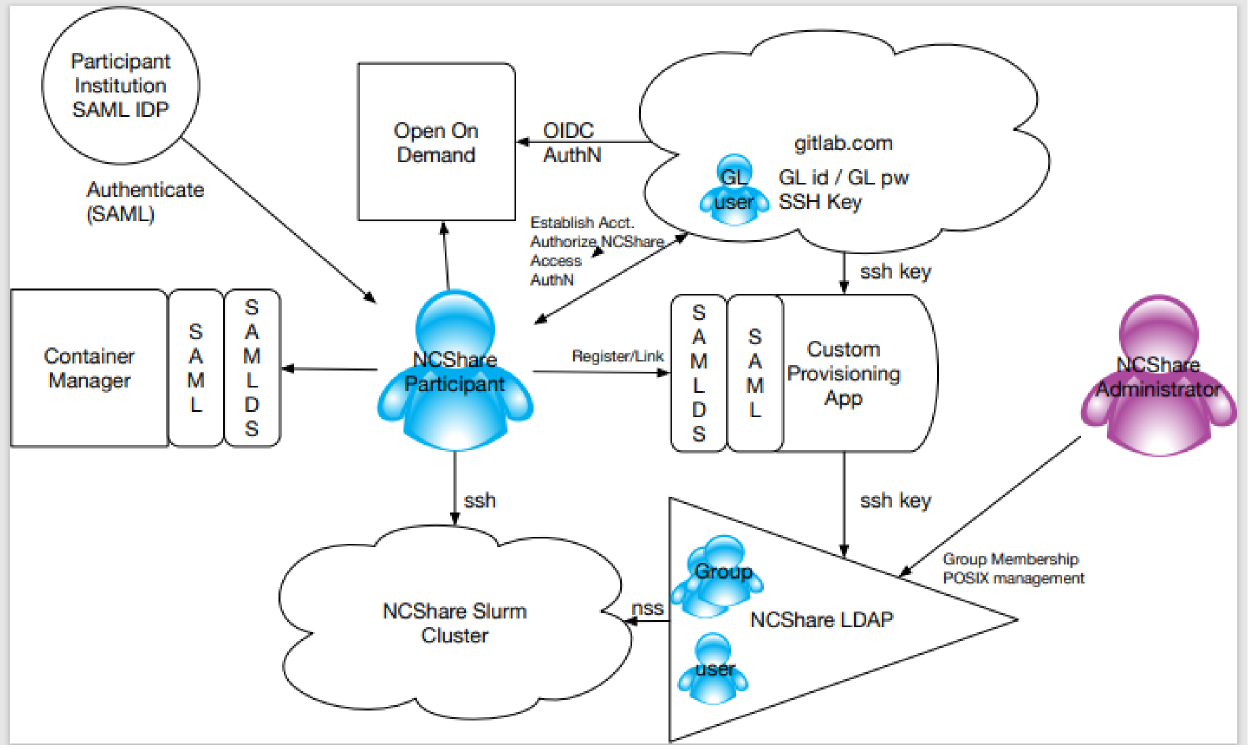

Figure 1 - Depicts the original IAM integration model used by NCShare and the rough institutional and user onboarding processes involved.

As participant institutions (particularly smaller, non-R1 schools within the state) have opted out of participation in the nationwide multilateral SAML federation (InCommon) and migrated their internal SSO solutions to commercial offerings from Microsoft, Okta, and Google and new participant institutions presented to NCShare, and as NCShare sought to expand into new compute offerings (including the anticipated addition of GPU compute services (GaaS) to NCShare's portfolio), the limitations of the initial IAM model became increasingly apparent. To address the shortcomings and friction points, NCShare chose to transition to a more robust IAM approach that models NCShare as a classical VO or Collaborative Organization and establishes first-class identities for NCShare participants linked to their institutional federated identities and managed through an NCShare COManage () instance. In this new environment, NCShare operates a COManage instance hosting a single CO (NCShare) that uses federated SAML identities to establish organizational anchor identities and automatically creates and provisions NCShare user identities based on those organizational identities into the NCShare LDAP, and a SAML-to-OIDC proxy (a Shibboleth IDP configured in proxy mode) that projects identities from the COManage-managed LDAP to web-based applications inside NCShare. The new model offers a number of improvements in manageability, flexibility, and end-user experience over the original model:

- GitLab.com dependency is eliminated. Users are no longer required to establish or use GitLab.com credentials to access any services inside NCShare (although some users may continue to use GitLab for other purposes).
- NCShare users now have first-class NCShare identities that are fully linked to their institutional identities via COManage, and can use their institutional SAML credentials to access all web-based applications inside NCShare
- Participating institutions need only maintain federated trust with two endpoints inside NCShare (the COManage instance and the SAML proxy), regardless of how many web-based services NCShare adds to its portfolio in future
- With first-class identities available for NCShare participants in COManage, it becomes possible to manage authorization groups, identity lifecycles, and other standard IAM objects and processes more effectively, and to delegate control of certain IAM features (such as authorization group memberships) to distributed institutional IT staff, freeing up NCShare staff time and increasing scalability across the entire project.

The transition required NCShare staff effort in a number of areas:

- Deploying and configuring a COManage instance for NCShare
- Customizing COManage PHP code to enable "one-click" federated identity linking and COManage identity provisioning – this greatly simplifies the onboarding process for new participants in NCShare – they use a single website to authenticate and get access without any additional work required by NCShare staff.
- Defining and configuring COManage enrollment flows for automated self-service enrollment, Linux cluster / POSIX identifier assignment, and manual (administrator-driven) enrollment
- Deploying a new openLDAP instance and configuring automatic provisioning from COManage to the LDAP instance
- Developing COManage API-based tools to implement annual identity attestation and automatic identity deactivation for COManage identities and provisioned services

- Establishing and implementing automated group management mechanisms inside COManage based on participants' institutional affiliations
- Establishing and implementing a delegated administration model to allow designated IT staff at participating institutions to manage memberships in COManage groups used to convey granular access to restricted data and computing resources inside NCShare (eg., GPU allocations in slurm)
- Deploying and configuring a Shibboleth IDP in SAML-to-OIDC and SAML-to-SAML proxy mode, with identity information sourced from the COManage-managed LDAP
- Negotiating InCommon registration for both the COManage instance and IDP proxy, and establishing R&S entity tags for both
- Negotiating bilateral trust relationships with each non-InCommon participant NCShare partner institution (currently 10 out of 22 NCshare institutions are bilaterally federated and 12 are federated through InCommon)
- Deploying and configuring a standalone SAML discovery service (based on RA21's SeamlessAccess tooling) for the NCShare pseudo-federation
- Reconfiguring Open On Demand to use the NCShare SAML-to-OIDC proxy for user authentication and reconfiguring Container Manager to use the SAML-to-SAML proxy for user authentication
- Manually migrating identity information and data associated with "legacy" entries in the original openLDAP directory to the new infrastructure (a one-time transition expense)

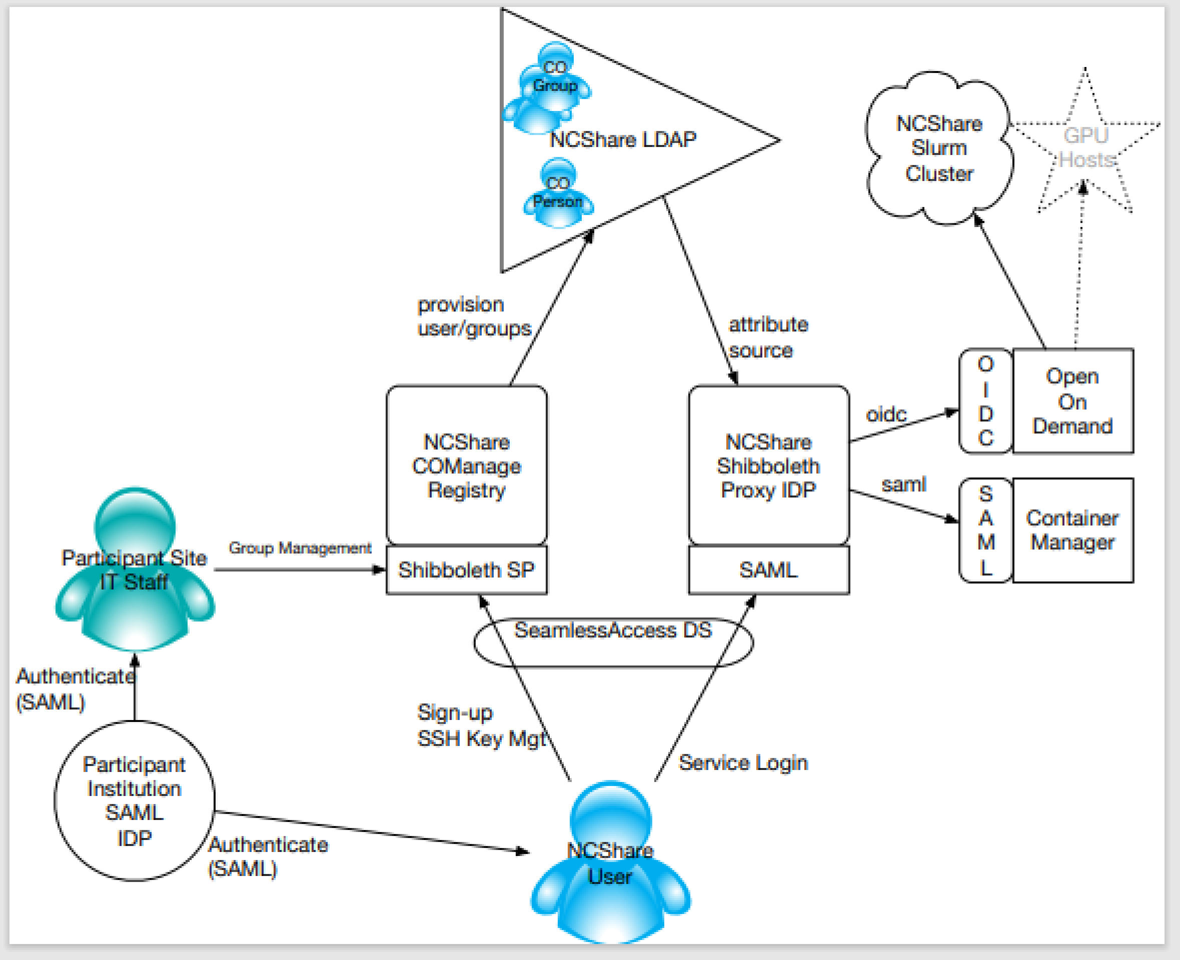

Figure 2 - Depicts the new COManage-driven infrastructure and associated user and administrator workflows.