To AD, or not to AD?

(that's NOT the question. yet.)

Reevaluating SDCC's AAI technical debt.

Matt Cowan, Jerome Lauret, Jason Smith, …

2024/05/16, SDCC TAB

Brief History

- Many years ago, for numerous technical reasons, SDCC deployed it's own AAI (Authentication and Authorization Infrastructure) instead of using BNL's AD (Active Directory)
  - Those reasons are (almost) gone (after AD changes for nsls2)
- Revisited using common AAI many times
  - only once since changes for nsls2, changes still too fresh then
- Currently running self supported FreeIPA 4.6.8-5.el7 on RHEL 7.9 (eol in ~2months, 2024/06/30, but planning ELS for RHEV, eol 2028/06/30)
- Support for IdM (RH tweaked version of FreeIPA) included with RHEL server/workstation licenses!

This *again*!?

Yup! well… sort of.

- Best practice to periodically revisit technical debt
- Especially when conditions change
- Little/no remaining IPA/Keycloak/Pidea expertise

But aiming different this time:

- Small incremental steps / bite size pieces to not choke
- Slowly pull in more people+cost+risk as we go
- Evaluate continuing at each step

Baseline considerations

- SDCC has several kinds of “accounts”
  - Generic user accounts
  - Service accounts - may be privileged within scope
  - Group accounts - may be privileged and accessed by a few people. ssh keys used to access #G#
  - Federated account - created near on the fly. #F#
  - Grid accounts
- On top of that, secondary groups are assigned
  - Request comes to be part of experiment X, SDCC validates with experiment
  - One account may belong to many groups / experiments
- Any solution must consider the points above
  - How do we overlay the “other” kind of accounts
  - Will we have access to group creation? Change group/user assignment a-la-NSLS?

Yet more considerations …

- MFA - how do we handle this?
  - NSLS2 uses AD + Duo, yubikeys if you cross a bastion host (going to an enclave). Easy to hand over one in their environment (most beamLine scientists come in person)
  - SDCC has over 2,000 active users, many users will NEVER come to BNL (but may only need AD+Duo)
  - Total number of users is fairly large (any solution should consider this and the associated cost)
  - Hybrid model work? Use cases are important to consider.
- Mobile users
  - A solution should consider that cell access may not be working (Duo like should support internet)
- Large scope
  - Many users requesting to change their password - would we direct them to itdhelp? What is group?
  - Would K5 still be used? How does it play with an AD based implementation?
  - Do we keep Emails separate or move this as well?

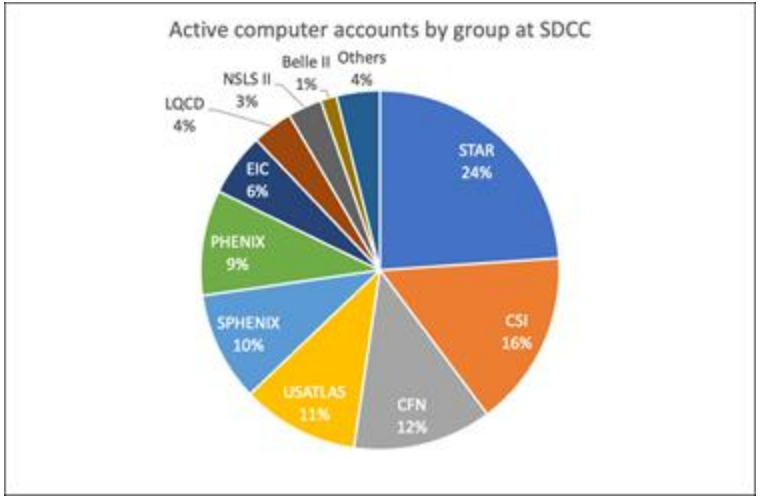

Brookhaven National Laboratory

Centrify? *(now called Delinea)*

- Centrify solution proven by nsls2 (~1.3k nodes?)
- SDCC = ~4k nodes!?
  - ~$200/node/year = ~$800k/year *(likely negotiable, but still prohibitive?)*
  - Impact on AD servers from adding that many direct clients?
  - Feasibility of migrating that many? As close to dropin replacement strongly desired
  - Failure behavior? BNL AD failure = 5 alarm fire, SDCC<->AD connectivity loss more realistic concern to try to engineer around/mitigate

An initial plan, proof of concept

- If generally agreed unified identity/auth/infra theoretically good…
- And no blatant barriers identified… *(use cases need to be carefully considered!)*

Propose:

- Doing presentation(s), so stakeholders aware & included, and gathering feedback *(have NOT considered all use cases yet!)*
- Start with a minimal feasibility test
  - Get minimal testbed working (2 IdM VMs in trust with AD, 3 client VMs: rhel[7-9])
  - Make it work with Duo
  - Make it work with ssh keys (any interest in sudo?)

Based on evaluation, proceed forward.

Duo

- Ability to use Duo is one of main benefits of AD
- RH supports IdM, including AD integration, but not Duo (or any indirect mfa through AD). RH pointed to examples of Duo working, but unsupported
  - pam_duo module is a possibility
- NSLS2 uses centrify - AD / Duo push integration exists
- We'd have support for IdM, including AD integration, just not Duo/mfa part.
- Currently self-support FreeIPA deployment.

Potential Future Steps

If feasibility test successful, potentially continue:

- Use agile methodology: reevaluate at each step, progressively take on more complexity+cost+risk!
- Implement mechanism enabling sdcc to create+manage AD groups (like nsls2)
- Broadly reach out to sdcc and stakeholders for use cases, expand testing
- Carefully architect migration plan
- Schedule and execute migration plan. Possible target window: soon after RHIC stops, before EIC work ramps up too much? Or maybe partially incremental with new independent experiments starting to use AD sooner?

Rough timeline

- pilot: aiming for next few to several months?
- intermediate steps: TBD in project planning
- migration target: Shortly after RHIC shutdown, before EIC rampup.
  Maybe possible to do incrementally, to some extent?

Questions?

more detailed view

Next steps:

- present to TAB and to liasons (below proceeds in parallel)
- organize, plan, and pursue pilot with smaller team: Dave, Stephen, James, Jason, myself, ???, others welcome
  - estimate resource requirements for pilot
  - evaluate
  - pilot planning
  - implement
  - test (rhel[7-9])
  - implement and test duo
  - implement and test ssh keys
  - iterate as needed

end-of-pilot

- reconvene larger team and reevaluate
- begin thorough project planning
- some subsequent steps (not comprehensive or in optimal order, just doc'ing some issues raised; reeval at each step)
  - broadly reach out within sdcc and stakeholders for use cases and add to test schedule
  - implement and test group mgmt (similarish to nsls2?)
  - research failure scenarios (focused on outages between AD and IdM as any AD issue would be 5 alarm fire), potential mitigations (caching possibilities?) and thouroughly test
  - evaluate options for potential progressive/incremental migration?
  - carefully architect a migration plan if all seems good to proceed
  - schedule and execute migration plan.

Supplemental info

- https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html/planning_identity_management/planning-integration-with-ad_planning-identity-management
- https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/9/html/installing_trust_between_idm_and_ad/index

Duo integration notes

… snippet from email thread from: kush@redhat.com, cc: mmiller@redhat.com …

Important:

IdM does not support OTP logins for Active Directory trust users.

However, I conducted some Google-fu and found some evidence of FreeIPA users getting Duo MFA to work through a Duo Authentication proxy server that would act as a RADIUS server. After a Duo Authentication proxy server is set up and accessible, these IdM steps to set up an external RADIUS connection should work in conjunction with Duo MFA.

- The documentation to configure a Duo Authentication proxy server with RADIUS and Duo-Only Secondary Authentication can be found here
  - Note that while Duo claims to have tested this to work on RHEL 7 and later, any issues related to the Duo proxy server will be viewed **as unsupported by Red Hat**

…