White Paper – MFA

Executive summary

Detta dokument beskriver ett förslag på MFA-lösning som en arbetsgrupp, bestående av representanter från samtliga lärosäten med SNIC-resurser, rekommenderar att användas vid autentisering till tjänster inom den nationella HPC-forskningsinfrastrukturen.

Att införa MFA kommer oundvikligen att kräva mer ansträngningar för användarna att få tillgång till tjänster. Genom att analysera olika praktiska användningsfall (se Appendix 1 - SNIC MFA Use Cases) har arbetsgruppen försökt identifiera balanserade avvägningar mellan användarvänlighet och behovet av att förbättra säkerheten och följa regulatoriska ramvillkor. Till exempel har fallet med användare som aktiverar HPC-jobb genom externt arbetsflödessystem analyserats. Vidare är vår rekommendation att autentiseringen till olika MFA-skyddade tjänster inom den nationella HPC-forskningsinfrastrukturen ska vara så lika som möjligt och undvika onödig mångfald som gör det krångligt för användarna.

Det finns krav, från bland andra MSB (MSB:s föreskrifter för statliga myndigheter; MSBFS 2020:6-8), att MFA måste användas när användare vill få åtkomst till information som har behov av extra skydd, som exempelvis känsliga personuppgifter eller sekretessbelagd information.

Ett önskemål från forskarna är att det ska se så lika ut som möjligt när de autentiserar sig till de olika tjänsterna. Det är dock av stor vikt att det extra steget vid autentisering, som MFA innebär, inte upplevs allt för omständigt, utan det ska snarare kännas tryggt att extra säkerhet används för att skydda informationen.

Arbetsgruppen rekommenderar att MFA implementeras även för autentisering till andra tjänster i den nationella forskningsinfrastrukturen (där MFA inte är ett obligatoriskt krav) för att användaren ska ha ett likartat sätt att autentisera sig på även där.

Några lärosäten inom SNIC-konsortiet har redan implementerat MFA och erfarenheter av genomfört installationsarbete och säkerhetskrav som ställts, ligger till grund för detta förslag.

Definitioner

IdP: Identity Provider. Identitetsintygare. Funktion som bekräftar att en användare av en tjänst eller ett nätverk med gemensam inloggning har rätt att logga in.

MFA: Multifaktorautentisering. I denna rapport använder vi beteckningen MFA för begreppen tvåfaktorsautentisering eller stark autentisering.

MSB: Myndigheten för Samhällsskydd och Beredskap.

MSBFS: MSB:s föreskrifter.

TOTP: Time-based One Time Password.

Avgränsningar

Detta dokument beskriver personliga inloggningar, inte maskin-maskin-inloggningar.

Nulägesanalys

Användarfall (use cases)

För att säkerställa att forskarna ska få ett så likartat sätt att autentisera sig på som möjligt på olika SNIC-resurser, har arbetsgruppen identifierat olika användarfall som finns när forskare/administratörer när använder/ administrerar tjänster inom den nationella HPC-forskningsinfrastrukturen.

De identifierade användarfallen finns beskrivna i Appendix 1 - SNIC MFA Use Cases.

I användarfallen har arbetsgruppen identifierat om forskningsdata innehåller information som kan anses vara i behov av utökat skydd (t.ex. känsliga personuppgifter, information som kan beläggas med sekretess eller företagshemligheter).

Enligt MSB:s föreskrifter (MSBFS 2020:6-7), krävs det att MFA används vid systemadministrativ åtkomst till informationssystem, och vid åtkomst till informationssystem som behandlar information som bedömts ha behov av utökat skydd.

Arbetsgruppen har bedömt att information som kan beläggas med sekretess och känsliga personuppgifter anses vara i behov av utökat skydd i form av multifaktorautentisering (MFA). I vissa användarfall när inte behov av utökat skydd anses behövas, rekommenderar vi ändå att det ska vara möjligt att använda MFA (för att det ska vara så likvärdiga sätt att autentisera sig på som möjligt).

Arbetsgruppen har även identifierat hur ofta forskarna autentiserar sig. Med den informationen har det föreslagits hur ofta forskarna behöver använda sig av MFA för att autentisera sig (varje gång, en gång per dag eller inte alls). Om en interaktiv session sträcker sig längre än rekommenderad period för hur ofta autentisering bör ske, bör det utredas för det specifika användarfallet om sessionen ska brytas eller inte.

Det finns även användarfall fall där autentisering sker i automatiserade processer, t.ex. arbetsbelastningshanteringssystem, där man inte kan tillhandahålla en TOTP varje gång.

I Appendix 2 - Required MFA finns en sammanställning över hur hög konfidentialitetsnivå informationen som hanteras i de olika användarfallen kan ha, om MFA är ett skallkrav eller en rekommendation, samt hur ofta MFA krävs vid autentisering.

Erfarenheter från existerande lösningar

En del tjänster inom den nationella HPC-forskningsinfrastrukturen har redan implementerat MFA och arbetsgruppen har dragit lärdomar från utfört implementationsarbete.

Arbetsgruppen har även tittat på vilka MFA-lösningar som internationella superdatorcenter har/planerar att implementera. Planen är att forskarna ska ha så likartade autentiseringssätt som möjligt till de flesta tjänster de använder.

Arbetsgruppen har även blickat ut i Europa för att se hur andra centra implementerat MFA. I Appendix 3 – White paper MFA – alternatives beskrivs tre olika flöden som kan vara bra att läsa.

Det ska dock tilläggas att icke-auktoriserade personer alltid försöker hitta nya sätt att attackera på och att det hela tiden sker en kontinuerlig utveckling av hur man bäst skyddar sig mot de nya attackerna. De lösningar som beskrivs i detta dokument kan i en snar framtid klassas som osäkra när nya angreppsmetoder dyker upp.

Teknisk lösning

Vi har identifierat krav som; att forskaren inte ska behöva skaffa sig ny utrustning (befintliga mobiler kan användas), att MFA-lösningen inte ska vara unik för enskild tjänst, att MFA-lösningen ska bygga på etablerad standard (inte kräva specifik produkt).

Existerande tekniska MFA-lösningar beskrivs i tre bilagor; Appendix 4 – TOTP IN SUPR, Appendix 5 – NSC-mfa och Bilaga 6 – Uppmax-mfa.

Under dessa förutsättningar, och med erfarenheter från existerande implementeringar, så bedömer arbetsgruppen att lösningen bör bygga på standarden Time-based One Time Password (TOTP), och att en autentiseringsapp av typen Authenticator (exempelvis Google Authenticator, Microsoft Authenticator eller Yubico Authenticator) installeras på användarens mobil.

Bilaga som beskriver tekniska lösningsförslag kopplade till de identifierade användarfallen; Bilaga 7 – MFA technical.

Användarsupport

Användarregistrering och initiering av MFA-lösningen är beroende på vilken IdP som används. För användare av den nationella HPC-forskningsinfrastrukturen som registreras via SUPR, bör MFA-supporten hanteras via SUPR-supporten.

Även rutiner för borttappade, eller vid byte av, mobiltelefoner och annan MFA-support bör kunna hanteras via SUPR-supporten.

Det bör dock utredas vidare hur support lämpligast ska skötas när annan IdP hanterar användarregistreringen.

Rekommendationer

Arbetsgruppen rekommenderar följande:

- MFA-lösningen ska bygga på standardlösning Time-based One Time Password (TOTP).
- En autentiseringsapp av typen Authenticator installeras på enhet separat från inloggande enhet, exempelvis användarens mobil.
- MFA rekommenderas att vara obligatorisk för autentisering till tjänster inom den nationella HPC-forskningsinfrastrukturen som har behov av utökat skydd.
    - MFA bör även erbjudas på övriga resurser inom den nationella HPC-forskningsinfrastrukturen. Se bilaga Required MFA för exempel.
- Användarregistrering bör ske via SUPR, eller annan likvärdig IdP.
- Användarsupport bör ske via SUPR-supporten, eller annan likvärdig IdP.
- Användarupplevelsen för nya MFA-installationer ska likna varandra.

Appendix 1 - SNIC MFA Use Cases

Services based on compute resources

HPC interactive access services based on general-purpose HPC resources

Connect to HPC system via SSH to start-up shell on login or visualization nodes

Use case name	UC01a: HPC SSH login Swedish researcher (at office or connected via VPN)
Primary actor	Swedish researcher
Stakeholders and interests	• *Researcher*: Wants to successfully start a shell on a login, visualization or similar node of a generally-purpose HPC system (in the following called target node)
Preconditions	• Researcher is registered in SUPR and does have an account on the HPC system • Researcher has an SSH client locally installed (OpenSSH, putty) • Researcher is at an office and maintains the same network connectivity the entire day (possibly due to use of VPN)
Success guarantee	Login-shell on target node started
Main success scenario	1. Researcher starts, if necessary, a VPN connection 2. Researcher start SSH client 3. Researcher provides all required credentials
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	The term “general-purpose HPC system” is used to distinguish this system from systems with enhanced security requirements, e.g. to support processing of sensitive data or industrial use cases.

Use case name	UC01b: HPC SSH login Swedish researcher (roaming)
Primary actor	Swedish researcher
Stakeholders and interests	• *Researcher*: Wants to successfully start a shell on a login, visualization or similar node of a generally available HPC system (in the following called target node)
Preconditions	• Researcher is registered in SUPR and does have an account on the HPC system • Researcher has an SSH client locally installed (OpenSSH, putty) • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	Login-shell on target node started
Main success scenario	1. Researcher starts SSH client 2. Researcher provides all required credentials
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	The term “generally available HPC system” is used to distinguish this system from systems with enhanced security requirements, e.g. to support processing of sensitive data or industrial use cases.

HPC interactive access services based on general-purpose HPC resources

Connect to HPC system via Thinlinc

Use case name	UC02: Remote desktop on HPC system (roaming)
Primary actor	Swedish researcher
Stakeholders and interests	• *Researcher*: Wants to successfully start a remote desktop
Preconditions	• Researcher is registered in SUPR and does have an account on the HPC system • Researcher has an ThinLinc client locally installed (see https://www.cendio.com/thinlinc/download) • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	Remote desktop is started
Main success scenario	1. Researcher starts SSH client 2. Researcher provides all required credentials
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	

HPC interactive access services based on general-purpose HPC resources

Connect to HPC system via JupyterHub service

Use case name	UC03: Access to JupyterHub service
Primary actor	Swedish researcher
Stakeholders and interests	• *Researcher*: Wants to successfully start a remote desktop
Preconditions	• Researcher is registered in SUPR and does have an account on the HPC system • Researcher has a web browser installed • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	A JupyterLab could be successfully started on target node (login node, compute node, viz node or similar)
Main success scenario	1. Researcher connects to JupyterHub through webbrowser 2. Researcher provides all required credentials 3. Researcher starts an existing or new JupyterLab
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	

Access to HPC resources dedicated to sensitive use cases (personal data, industrial use)

Use case name	UC04: HPC SSH login Swedish researcher with access to sensitive data (roaming)
Primary actor	Swedish researcher
Stakeholders and interests	• *Researcher*: Wants to successfully start a shell on a login, visualization or similar node of a generally available HPC system (in the following called target node) with access to sensitive data as well as the option to use an SSH connection for transferring sensitive data into the system or back
Preconditions	• Researcher is registered in SUPR and does have an account on the HPC system • Researcher has an SSH client locally installed (OpenSSH, putty) • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	Login-shell on target node started and sensitive data accessible
Main success scenario	1. Researcher starts SSH client 2. Researcher provides all required credentials
Special requirements	Totally separated projects using virtual cluster on openstack (meaning they can see no users, data, processes, job queues or anything from another sensitive project). No internet access from the inside to the outside. To be able to control all data leaving the system. This makes it a bit cumbersome to work with, but that is for minimizing the risk of data leakage (by accident, carelessness or even misconduct). So to clarify it even more: You can't reach any IP-addresses outside its own project.
Frequency of occurrence	Multiple times per day
Example	bianca

Science gateway service (e.g. AiiDA) connecting to HPC system for launching jobs on general-purpose HPC resources

Use case name	UC05a: Launching HPC jobs through a Science Gateway
Primary actor	User (anywhere in the world, not known to the data centre)
Stakeholders and interests	• *User*: Wants to start a pre-installed workflow on an HPC system through a web-portal service
Preconditions	• User has access to the science gateway service through an identity that may not be known to the HPC system operator and possibly a weak authentication mechanism • Science gateway has necessary credentials for starting jobs on HPC system (e.g. via a so-called service account) • Pre-defined workflows with pre-installed HPC applications are available, i.e. the User is not able to start any other HPC application • Availability of a *Service Account* that is used for executing the workflow on the HPC system
Success guarantee	Workflow is successfully started
Main success scenario	1. User connects to Science Gateway 2. User configures a workflow via the Science Gateway 3. User starts workflow from the Science Gateway for execution on the HPC system
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	The science gateway is provided by a portal service provider on basis of an SLA involving the e-infrastructure services provider, which operates the HPC system.

External workflow system launching jobs on general-purpose HPC resources

Use case name	UC05b: Launching HPC jobs through external workflow system
Primary actor	Swedish researcher
Stakeholders and interests	• *Researcher*: Runs workflow tool on external system that established an SSH connection to the HPC system for submitting jobs
Preconditions	• Researcher is registered in SUPR and does have an account on the HPC system • Researcher has an SSH client locally installed that can be invoked by the workflow tool.
Success guarantee	The workflow tool completes all necessary job submissions
Main success scenario	1. User provides the necessary credentials to the workflow system 2. User starts workflow tool that connects to HPC system for interacting with Slurm
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	Example workflow framework: httk (https://github.com/httk/httk.git)

Access to private cloud management portal (e.g. OpenStack dashboard)

Use case name	UC06: Science cloud
Primary actor	Swedish researcher
Stakeholders and interests	• *Researcher*: Wants to start an infrastructure service, e.g. a virtual server, through the management portal, e.g. an OpenStack Dashboard
Preconditions	• Researcher is registered in SUPR and does have an identity from an IdP supported by the cloud system • Researcher has a web browser locally installed • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	Connection to the cloud management portal is successful and a service can be started
Main success scenario	1. Researcher connects to management portal through local web browser 2. Researcher provides all required credentials
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	

Administrator access to HPC management nodes and services

Use case name	UC07: HPC SSH login
Primary actor	System Administrator
Stakeholders and interests	• *System Administrator*: Wants to access management resources for monitoring and configuration • *Security Officer*: Wants to ensure that security protocols (SNIC/MSBFS) are followed
Preconditions	• System Administrator does have a local account on the relevant systems • System Administrator has an SSH client locally installed (OpenSSH, putty) and a public SSH key installed on the target system • System Administrator connects from a network that is considered secure
Success guarantee	System Administrator can access management services
Main success scenario	1. System Administrator starts, if necessary, a VPN connection 2. System Administrator starts SSH client 3. System Administrator provides all required credentials
Special requirements	MSBFS requires MFA for system administrators
Frequency of occurrence	Multiple times per day
Miscellaneous	

Services based on storage resources

Interactive access service to data management nodes with access to storage resource connected to general-purpose HPC resource (i.e. PFS) using SSH

Use case name	UC11: Data management access via SSH
Primary actor	Swedish Researcher
Stakeholders and interests	• *Researcher*: Wants to have access to a data management node from where she/he has access to her/his data stored in a PFS, which is attached to the HPC system, and can move data through the external internet
Preconditions	• Researcher is registered in SUPR and does have an account on the HPC system • Researcher has an SSH client locally installed (OpenSSH, putty) • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	Login-shell on target node started
Main success scenario	1. Researcher starts SSH client 2. Researcher provides all necessary credentials
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	The term “general-purpose HPC resource” is used to distinguish this system from systems with enhanced security requirements, e.g. to support processing of sensitive data or industrial use cases. This implicitly means that in this use case does not foresee access to sensitive data.

Authenticated write access to storage resources via S3 or Swift

Use case name	UC12a: Authenticated write access to a general-purpose object storage via S3/Swift
Primary actor	Any User (not necessarily based in Sweden) with an identity from an accepted IdP
Stakeholders and interests	• User wants to upload objects to an existing bucket from different network locations (including office and home office)
Preconditions	• S3/Swift client locally installed (e.g. OpenStack Swift python client, Cyberduck)
Success guarantee	Object upload is successfully started
Main success scenario	1. User obtains OIDC or SAML token by providing all necessary credentials 2. User starts client for connecting to the S3/Swift API
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	The term “general-purpose object storage” indicates that this storage is not meant to hold sensitive data.

Authenticated read access to storage resources via S3 or Swift

Use case name	UC12b: Authenticated read access to a general-purpose object storage via S3/Swift
Primary actor	Any User (not necessarily based in Sweden) with an identity from an accepted IdP
Stakeholders and interests	• User wants to download existing objects from different network locations (including office and home office)
Preconditions	• S3/Swift client locally installed (e.g. OpenStack Swift python client, Cyberduck) • User has an identity from an accepted IdP
Success guarantee	Object download is successfully started
Main success scenario	1. User obtains OIDC or SAML token by providing all necessary credentials 2. User starts client for connecting to the S3/Swift AP
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	The term “general-purpose object storage” indicates that this storage is not meant to hold sensitive data.

Direct access to storage resources via WebDAV

Use case name	UC13: Authenticated data transfer via WebDAV
Primary actor	Swedish Researcher
Stakeholders and interests	• Researcher wants to upload file • Researcher wants to download file
Preconditions	• Researcher is registered in SUPR and has access to a SNIC storage project • If using client certificate for authentication, researcher has registered this certificate in SUPR • WebDAV client or web browser installed locally
Success guarantee	File upload and download are successful
Main success scenario	1. Researcher can authenticate to WebDAV server 2. Researcher can transfer files
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	Most WebDAV clients only support password and client certificate authentication.

Direct access to storage resources via gridftp, SRM, xRootd, dcap. sftp

Use case name	UC14: Authenticated data transfer using gridftp/SRM/xRootd/dcap/sftp
Primary actor	Swedish researcher
Stakeholders and interests	• Researcher wants to upload file • Researcher wants to download file
Preconditions	• Researcher has an X.509 certificate • Researcher has relevant clients installed locally • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	File upload and download is successful
Main success scenario	1. Researcher obtains a proxy certificate 2. Researcher starts file upload or download
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	

Direct access to storage resources via NFS

Use case name	UC15: Authenticated data transfer via NFS
Primary actor	Swedish Researcher
Stakeholders and interests	• Researcher wants to upload file • Researcher wants to download file
Preconditions	• Researcher has an NFS client • Researcher has an access token that the NFS client can use
Success guarantee	File upload and download are successful
Main success scenario	1. Researcher obtains an access token 2. Researcher can mount the NFS share 3. Researcher can read and write files via NFS
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	In practice authenticated NFS usage requires Kerberos since that is what implementations support today. In this Use Case no access to sensitive data is assumed.

Direct access to storage resources via SFTP

Use case name	UC16: Authenticated data transfer via automated SFTP
Primary actor	Swedish Researcher
Stakeholders and interests	• Researcher wants to upload file • Researcher wants to download file
Preconditions	• Researcher has an SFTP client • Researcher has an access token that the SFTP client can use • No user interaction can be required after obtaining the access token
Success guarantee	File upload and download are successful
Main success scenario	1. Researcher obtains an access token 2. Researcher stores the access token and makes it available to the SFTP client 3. Automated jobs scheduled by the researcher can read and write files via SFTP
Special requirements	No user interaction is possible when the SFTP client authenticates
Frequency of occurrence	Multiple times per day
Miscellaneous	

Direct access to storage resources via SFTP

Use case name	UC17: Authenticated data transfer via interactive SFTP
Primary actor	Swedish Researcher
Stakeholders and interests	• Researcher wants to upload file • Researcher wants to download file
Preconditions	• Researcher has an SFTP client • Researcher has credentials providing access to SFTP server
Success guarantee	File upload and download are successful
Main success scenario	1. Researcher starts SFTP client 2. Researcher provides all necessary credentials 3. Researcher can upload and/or download files
Special requirements	
Frequency of occurrence	Multiple times per day
Miscellaneous	

Other services

User connects to SUPR to change information or access non-public information

User attributes not related to authentication

Use case name	UC21: User access to SUPR for changing user attributes not related to authentication
Primary actor	Swedish researcher
Stakeholders and interests	• Researcher wants to connect to the SUPR portal to update his profile and change relevant attributes (including attributes like SSH public keys that may affect authentication to other services)
Preconditions	• Researcher has an account on SUPR • Researcher has web browser installed locally • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	Researcher can connect to SUPR and perform the targeted changes
Main success scenario	1. Researcher starts local browser 2. Researcher connects to SUPR and provides all necessary credentials
Special requirements	
Frequency of occurrence	Once a week or less often
Miscellaneous	

User connects to SUPR to change information or access non-public information

User attributes related to authentication credentials (e.g. public SSH keys)

Use case name	UC22: User access to SUPR for changing user attributes related to authentication
Primary actor	Swedish researcher
Stakeholders and interests	• Researcher wants to connect to the SUPR portal to change relevant attributes that may affect authentication to other services (e.g., SSH public keys)
Preconditions	• Researcher has an account on SUPR • Researcher has web browser installed locally • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	Researcher can connect to SUPR and perform the targeted changes
Main success scenario	1. Researcher starts local browser 2. Researcher connects to SUPR and provides all necessary credentials
Special requirements	
Frequency of occurrence	Once a week or less often
Miscellaneous	

PI connects to SUPR to change information related to projects

Use case name	UC23: Update of project information in SUPR
Primary actor	Swedish Researcher
Stakeholders and interests	• Researcher wants to connect to the SUPR portal to change information related to projects where she/he acts as PI
Preconditions	• Researcher has an account on SUPR • Researcher has web browser installed locally • Researcher is roaming between multiple locations/networks and does not have a stable IP address
Success guarantee	Researcher can connect to SUPR and perform the targeted changes
Main success scenario	1. Researcher starts local browser 2. Researcher connects to SUPR and provides all necessary credentials
Special requirements	
Frequency of occurrence	Once a week or less often
Miscellaneous	

Direct access to SENSE data storage via sftp (similar to UC14)

Use case name	UC24: Sens data access sftp
Primary actor	Swedish researcher
Stakeholders and interests	*Researcher*: Wants to successfully download sensitive data C3
Preconditions	Researcher has an SSH client locally installed (OpenSSH, putty)
Success guarantee	Files with correct checksums
Main success scenario	Downloaded and uploaded files with correct data
Special requirements	Detailed logging of transactions, totally separated projects (se definition in UC4)
Frequency of occurrence	Multiple times per day
Examples	Grus, bianca-sftp

HPC interactive access services for SENSE HPC resources Connect to HPC system via web based ThinLinc. (similar to UC04)

Use case name	UC25: Sens web based ThinLinc
Primary actor	Swedish researcher
Stakeholders and interests	*Researcher*: Wants to use HPC resources in graphical environment
Preconditions	Researcher has a supported browser installed (Chrome/Mozilla)
Success guarantee	Remote desktop is started
Main success scenario	
Special requirements	Totally separated projects, no internet (see UC4 for definition)
Frequency of occurrence	Multiple times per day
Examples	bianca

Sensitive data query/aggregation via web-based service

Use case name	UC26: Sensitive data query/aggregation via web-based service
Primary actor	Any User (not necessarily based in Sweden) with an identity from an accepted IdP
Stakeholders and interests	User wants to initiate aggregation of or query on sensitive data from different network locations (including office and home office)
Preconditions	User has an identity from an accepted IdP
Success guarantee	Query/data aggregation operation is initiated
Main success scenario	1. User obtains OIDC or SAML token by providing all necessary credentials 2. User connects via a web-browser to the service and starts data query/aggregation operation 3. Service response is returned in browser
Special requirements	Web-based service returns data that is non-personal according to GDPR, i.e. while the data query/aggregation operation is performed on personal data accessible to the service, only data is returned that does not allow re-identification of data subjects.
Frequency of occurrence	Multiple times per day
Miscellaneous	Such service was envisioned by the Human Brain Project's Medical Information Platform (https://www.humanbrainproject.eu/en/medicine/medical-informatics-platform/overview/)

Appendix 2 - Required MFA

Use case	Confid. level	2FA mandatory?	2nd factor frequency (minimum)	Comment
UC01a: HPC SSH login Swedish researcher (at office or connected via VPN)	1	mandatory	1/day	Access to dual-use technology but no access to sensitive data
UC01b: HPC SSH login Swedish researcher (roaming)	1	mandatory	1/day	Access to dual-use technology but no access to sensitive data
UC02: Remote desktop on HPC system (roaming)	1	mandatory	1/day	Access to dual-use technology but no access to sensitive data
UC03: Access to JupyterHub service	1	mandatory	1/day	Access to dual-use technology but no access to sensitive data
UC04: HPC SSH login Swedish researcher with access to sensitive data (roaming)	3	mandatory	always	Access to sensitive data, legal requirement
UC05a: Launching HPC jobs through a Science Gateway	1	never	n/a	Access to dual-use technology but no access to sensitive data
UC05b: Launching HPC jobs through external workflow system	1	mandatory	1/day	Access to dual-use technology but no access to sensitive data
UC06: Science cloud	1	mandatory	1/day	Control on infrastructure
UC07: HPC SSH login (system adminsitrator)	3	mandatory	always	Access to sensitive data; privileged control on infrastructure
UC11: Data management access via SSH	1	optional	1/day	No access to sensitive data
UC12a: Authenticated write access to a general-purpose object storage via S3/Swift	1	never	n/a	No access to sensitive data
UC12b: Authenticated read access to a general-purpose object storage via S3/Swift	1	never	n/a	No access to sensitive data
UC13: Authenticated data transfer via WebDAV	1	never	n/a	No access to sensitive data
UC14: Authenticated data transfer using gridftp/SRM/xRootd/dcap/sftp	1	never	n/a	No access to sensitive data
UC15: Authenticated data transfer via NFS	1	never	n/a	No access to sensitive data
UC16: Authenticated data transfer via automated SFTP	1	never	n/a	No access to sensitive data
UC17: Authenticated data transfer via interactive SFTP	1	never	n/a	No access to sensitive data

UC21: User access to SUPR for changing user attributes not related to authentication	3	mandatory	always	Access to sensitive data
UC22: User access to SUPR for changing user attributes related to authentication	3	mandatory	always	Access to security relevant data
UC23: Update of project information in SUPR	2	never	n/a	No access to sensitive data
UC24: Sens data access sftp	3	mandatory	always	Access to sensitive data
UC25: Sens web based ThinLinc	3	mandatory	always	Access to sensitive data
UC26: Sensitive data query/aggregation via web-based service	1	never	n/a	Service per definition returns non-sensitive data
UC01a: HPC SSH login Swedish researcher (at office or connected via VPN)	1	mandatory	1/day	Access to dual-use technology but no access to sensitive data
UC01b: HPC SSH login Swedish researcher (roaming)	1	mandatory	1/day	Access to dual-use technology but no access to sensitive data

Appendix 3 – White paper MFA – alternatives

Alternative Authentication Flows

At different sites throughout Europe different authentication flows involving multi-factor authentication are explored. Monitoring such efforts is important as many Swedish HPC users are also using other facilities in Europe. A certain uniformity of the chosen authentication flows can help to reduce the additional complexity for users. Furthermore, some of the considered authentication flows are more suitable in for a distributed HPC infrastructure where services provided by different sites are federated. The focus in the following on HPC-based services for which SSH-based connections are used.

Flow A

Periodic provisioning of a second authentication factor can be enforced for a setup where SSH key-pair based authentication is used by mandating periodic enabling of the public SSH key that is stored on the service, to which the user wants to connect. A schematic overview of the authentication is shown in this figure:

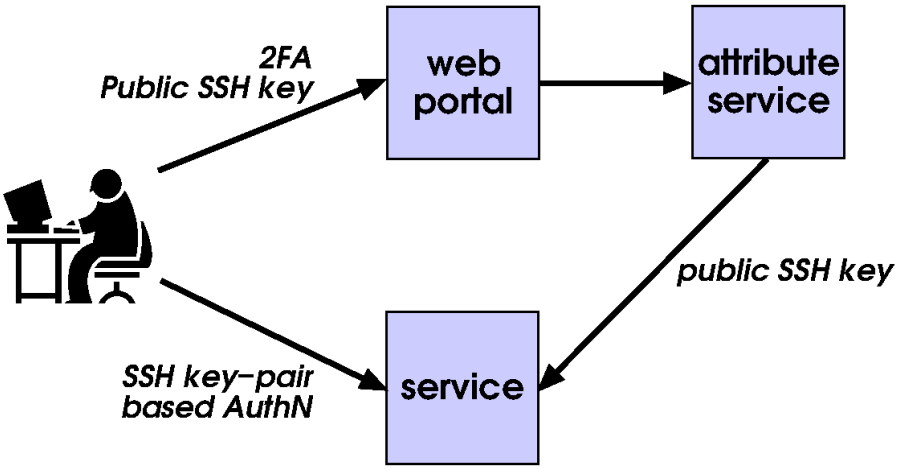

The following steps are foreseen:

1. The user connects to a web portal and authenticates with two factors, namely a password a one-time-password (OTP).
2. The user provides or re-activates its public SSH key, which is stored in an attribute service.
3. The service periodically deletes public SSH keys that are considered inactive and the attribute service pushes active SSH keys to the service.
4. The user connects to the service using its local private and the remote public SSH key.

Flow B

Provisioning of a second factor for SSH key-pair based authentication can also be integrated by foreseeing temporary SSH keys. Again, the user authenticates with a web portal using two factors. The portal will generate an SSH key-pair, which can only be used for a limited period of time. The authentication flow is shown in the following figure:

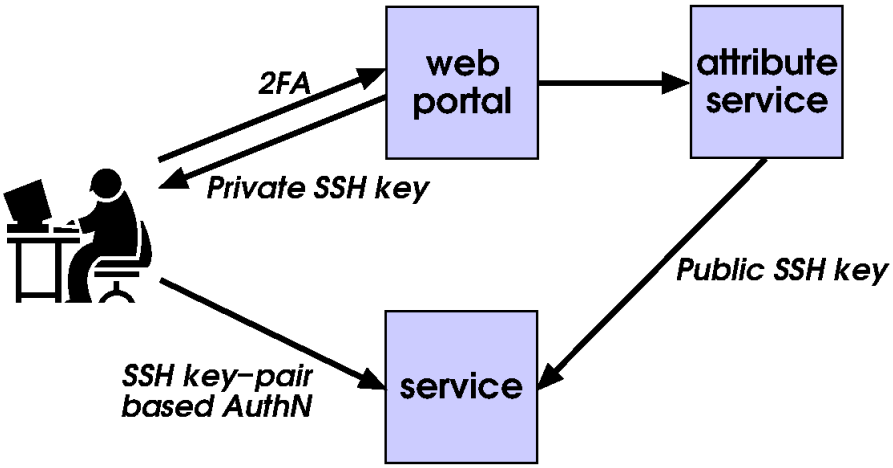

Flow C

The architecture outlined in this subsection assumes the availability of an IdP Proxy service, which is foreseen in federated infrastructures that are based on the AARC blueprint architectures[1]. The authentication architecture supports multiple authentication flows and, therefore, provides the flexibility to adapt to different site policies. It would, e.g. support an authentication flow, which is very similar to flow A. Furthermore, it will support authentication using temporary certificates. A schematic overview is provided in the following figure:

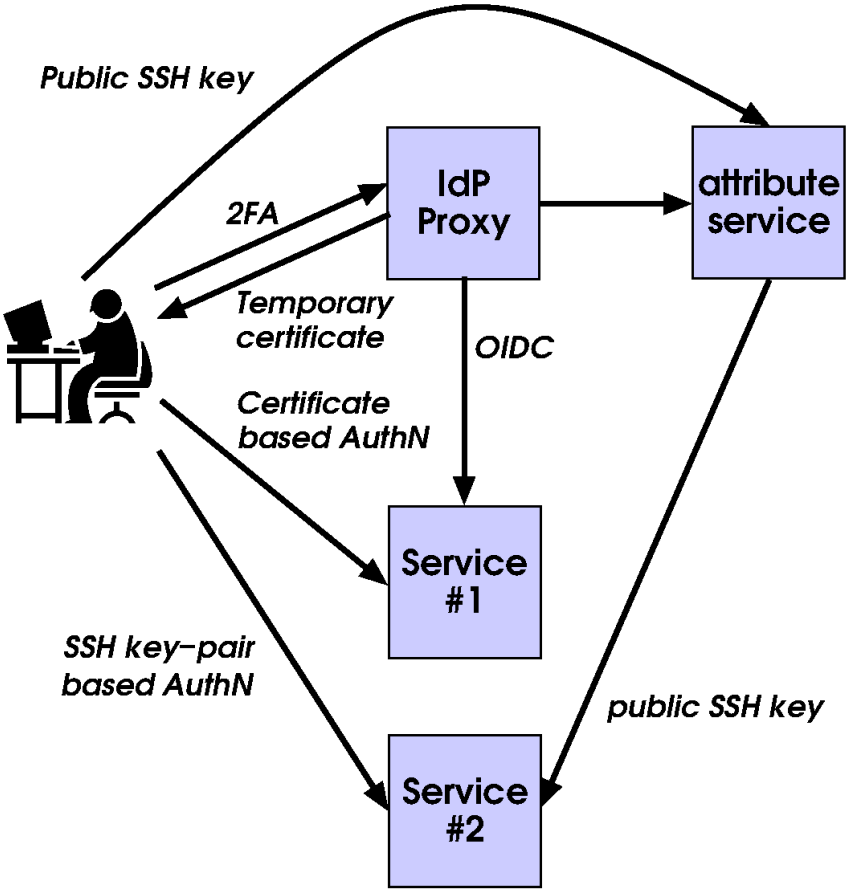

This architecture is currently envisaged for implementation by the Fenix HPC centres using Satosa as technology for realising a Proxy IdP. Satosa is also used for MyAccessID[2], but does as of today not yet support 2FA.

1 C. Kanellopoulos, C. et al., "AARC Blueprint Architectures", Tech. rep., AARC project (2017), https://aarc-project.eu/wp-content/uploads/2017/05/DJRA1.2-AARC-Blueprint-Architectures-1.pdf

2 https://wiki.geant.org/display/MyAccessID/MyAccessID+Home

Appendix 4 – TOTP IN SUPR

As currently deployed

SUPR uses TOTP as the second factor for authentication when that is enabled. This is on top of and independent of the primary authentication, that may be 1) federated login via SWAMID, 2) email and password (local to SUPR) or 3) client certificate (registered in SUPR). SUPR asks for a valid verification code from the TOTP app after the primary authentication has succeeded.

Enabling two-factor authentication via TOTP is mandatory for users with Staff/Administrator privileges and for PIs and proxies in SNIC SENS projects, and NGI Delivery projects (if marked as containing sensitive data). If such users have not enabled two-factor authentication, they are forced to do so when logging in. Other users may voluntarily enable two-factor authentication.

Enabling TOTP is done using the standard method: a page presents information about TOTP apps and shows the secret as a QR code (suitable for almost all users) and as text (for the few users who cannot scan the QR code). The user then gets the secret into the TOTP app and confirms this by entering a first generated verification code to confirm. At the same time, the user is encouraged to verify that their phone numbers in SUPR are correct.

Users who lose their TOTP app or cannot get it to work will contact SUPR support to have the TOTP removed so they can login without it (and enable it again at once if required). SUPR staff verifies the request, typically by calling a registered phone number in SUPR to talk to the user.

Developed but currently not in production

We have developed a function where a user can request a recovery code that is then written down and stored in a secure place. That code can be used a single time to remove (if allowed) or replace the TOTP secret.

Appendix 5 – NSC-mfa

NSC functional requirements for MFA implementation

Focused on interactive SSH logins to compute cluster login nodes

- Defend against stolen user credentials
- File system access when logged in to the cluster is not reliant on the SSH login credential
- Accessing other storage system such as Swestore is out of scope
- Each cluster is its own MFA domain to reduce dependencies and shared secrets

No requirements for hardware a normal user does not already have access to.

- We had already tried Yubikeys (OTP version) for Swestore and it had not been a success
- Only realistic hardware option left is mobile phones

Client implementation at no cost.

- TOTP deemed to be most suitable option
    - Multiple free implementations
    - Supports multiple clusters in one app
    - Already used by others (like SUPR)

Must integrate well with NSC account system.

- The NSC account and project system/database (and user portal) could be extended to support TOTP secrets

Solution must not be SNIC specific.

- NSC runs clusters for other organisations as well
- Some user overlap with SNIC systems, so the same technology should be used for user familiarity

Automated file transfers should still be possible.

- Support automated logins not using MFA
- Logins not using MFA may be restricted

System admin logins are not the focus and may use another mechanism.

- Stricter requirements for administrative access due to MSBFS regulations

Bilaga 6 – Uppmax-mfa

Genomgång av befintlig (och tidigare använd) MFA på UPPMAX

Admins (systemexperter – inte applikationsexperter)

Tidigare användes en proprietär SMS-lösning (Duo).

Nu används yubikey eller TOTP, den senare är att betrakta som nödlösning.

Användare

Olika varianter med TOTP har använts.

System för ickekänslig data

Rackham har haft nån form av TOTP tvåfaktor sen 2017. Bara krävd om anslutande IP anses komma utanför Sverige eller utanför SUNET. Implementerad med separat inmatning medelst "forced command", vilken går tämligen enkelt för användaren att lirka sig förbi.

System för känslig data

Irma har sen start 2015 haft TOTP med vissa särskilda undantag. Implementeringen liknar Rackhams men har också en implemeterad "gracetime" som gör att man inte behöver ge TOTP för varje access.

Bianca har sen start 2016 haft TOTP implementerad på centrala inloggningsservern där TOTP-koden matas in direkt efter det vanliga lösenordet i "lösenordsfältet". Denna har gällt för alla utan undantag, även för Biancas system för filöverföring med sftp ("Wharfen").

*Grus (*ett system enbart avsett för filöverföringar. TOTP implementerad på likartat sätt som i Bianca.

Miarka. TOTP aktiverad sen start på både nyckel- och lösenordsinloggning. implementerad med pam-modul och sshdconfig "keyboard-interactive". Kort sagt: som ssh är tänkt att haka på denna typ av funktionalitet.

UPPMAX TOTP-autenticeringsserver

Implementation

En egenskriven TOTP-server i python som en CGI-applikation till lighttpd. Denna svarar på HTTPS antingen med GET (våra gamla lösningar) eller med POST (Miarka).

UPPMAX implementation stödjer flera accesser för samma TOTP-kod inom tidsintervallet., vilket kan anses tveksamt. För att thinlinc ska fungera så krävs att det fungerar åtminstone två gånger med samma kod.

Denna server saknar i dagsläget redundans.

Registrera TOTP på UPPMAX

Man dirigeras till en SUPR-inloggning där man får bekräfta vem man är och sen dirigeras man till en

UPPMAX-server där man använder QR-kod "på normalt vis" för TOTP-registrering i app.

Denna server gör hemligheten tillgänglig under ca en halvtimme så autenticeringsservern kan hämta den, sen raderas den.

Denna server saknar i dag redundans.

Önskade förbättringar på UPPMAX

Allra helst skulle UPPMAX vilja ha en gemensam hemlighet för alla SNIC-center och SUPR, men då detta känns som ett väldigt stort steg att ta (både politiskt och tekniskt) så är det snarast steget “Registrera TOTP på UPPMAX” som vi skulle vilja hanterades i SUPR och på enhetligt sätt med de övriga centren. Med en sån lösning så är det ganska tydligt och klart att man kan säga att SNIC har en gemensam MFA-lösning, även om centrena gör vissa egna val i huruvida man ska ha olika hemligheter för olika kluster (NSCs approach) eller på hela centret (UPPMAX approach).

Denna typ av ändring är även SUPR-folket positiva till.

Bilaga 7 – MFA technical

Grundkoncept för interaktiv användarinloggning

TOTP: Time-based One Time Password

PAM: Pluggable Authentication Module

MFA-lösningen bygger på TOTP-app på separat enhet (i normalfallet en mobiltelefon) samt konfiguration av sshd och ett par PAM-moduler. Syftet är att använda ssh:s tänkta mekanismer för "keyboard-interactive" -- inga specialvarianter med force-command eller varianter där lösenordet och TOTP-koden mixas.

Lösningen för att slippa ge TOTP-koden allt för ofta är att konfigurera ssh-klienten med control/master (ssh-multiplexing).

Hur TOTP-hemligheten lagras är inget som påverkar användarupplevelsen, och är därför inget som behöver bestämmas gemensamt.

Varje användare måste sätta upp en TOTP-hemlighet för sitt konto. Då de flesta TOTP-appar har stöd för att antingen skriva in en lång kod eller med mobilkamera läsa en QR-kod görs detta lämpligen via inloggning på en webbsida alternativt med utskick via post. Smidigast för användare både vid nya konton och återställning av TOTP-hemlighet är om autentisering kan ske via webbsida. Här kan SWAMID används för majoriteten av användare, men alternativ måste finnas för att täcka in samtliga användare. För SNIC-system bör detta hanteras i SUPR.

Support för TOTP utgörs efter driftsättning (inkl. att få igång alla befintliga användare) huvudsakligen av återställning av TOTP-hemligheter. Byten av mobiltelefoner sker relativt ofta, och det är inte alla användare som har möjlighet att kopiera sin TOTP-hemlighet vid ett sådant byte. Rutinen för detta måste alltså både finnas på plats och vara rimligt enkel.

Hur många system som ska dela på en TOTP-hemlighet är en avvägning mellan fler faktorer utan uppenbart svar. Att dela hemlighet mellan flera huvudmän bör undvikas då det bör hållas inom en säkerhetsorganisation.

Implementering av UC01a (HPC SSH login Swedish researcher (at office or connected via VPN))

Normal SSH-klient kombinerad med mobiltelefon+TOTP-app för användaren. Här kan ControlMaster i SSH användas för att användaren har en stabil uppkoppling under dagen och då bara behöver autentisera via TOTP en gång.

Exempel på hur det ser ut för en användare (första inloggning, lösenord + TOTP används):

```
myuser@mylaptop:~$ ssh snicuser@cluster.snicsite.se
Password: <==== ENTER YOUR NORMAL PASSWORD HERE
Verification code: <===== ENTER THE 6-DIGIT CODE FROM YOUR APP HERE
```

Användare utan mobiltelefon kan använda en hårdvarunyckel (t.ex. Yubikey) kombinerad med en applikation på datorn för att få tillgång till TOTP.

Implementering av UC01b (HPC SSH login Swedish researcher (roaming))

Är i praktiken samma som UC01a, förutom att ControlMaster inte kan användas om IP-adressen ändras och/eller sessioner bryts på grund av datorn går ner I viloläge. Detta kan då leda till att TOTP måste anges av användaren för varje inloggning.

Implementering av UC02 (Remote desktop on HPC system (roaming))

Om Thinlinc används så används SSH för kommunikationen och serversidan är den samma som för "HPC SSH login". Dock så har klienten för Thinlinc vissa begränsningar som gör att endast lösenord kan användas om även TOTP krävs. För att använda SSH-nyckel+TOTP måste modifieringar av Thinlinc-klienten göras.

Implementering av UC03 (Access to JupyterHub service)

Implementering av UC04 (HPC SSH login Swedish researcher with access to sensitive data **(roaming))**

Går att göra exakt som UC1.

Implementering av UC05 (Launching HPC jobs through a Science Gateway)

Implementering av UC06 (Science cloud)

Implementering av UC07 (*HPC SSH login)*

Här kan antingen samma lösning som för UC01a användas, eller så kan hårdvarunycklar (t.ex. Yubikey) krävas eftersom det kan köpas in till samtliga anställda. Att tillhandahålla hårdvarunycklar till alla användare får anses vara orealistiskt.

Implementering av UC11 (Data management access via SSH)

Går att implementera på samma sätt som UC01.

Implementering av UC12a (Authenticated write access to a general-purpose object storage via S3/Swift)

Implementering av UC12b (Authenticated read access to a general-purpose object storage via S3/Swift)

Implementering av UC13 (Authenticated data transfer via WebDAV)

De klienter som används har stöd för antingen autentisering via X.509-certifikat eller användarnamn/lösenord. den mån känsliga data förekommer kan man tänka sig att kräva att smarta kort används för lagring av certifikat, men annars finns begränsade möjligheter att införa MFA.

Implementering av UC14 (Authenticated data transfer using gridftp/SRM/xRootd/dcap)

I detta fall används huvudsakligen X.509-certifikat för autentisering. I den mån känsliga data förekommer kan man tänka sig att kräva att smarta kort används för lagring av certifikat, men annars finns begränsade möjligheter att införa MFA.

Implementering av UC15 (Authenticated data transfer via NFS)

För Kerberos kan det finnas möjlighet att integrera en andra faktor med hjälp av preauth-mekanismen. Då de biljetter som utfärdas av en KDC har begränsad livslängd kan det utnyttjas för att konfigurera önskad frekvens av förnyad autentisering.

Implementering av UC16 (Authenticated data transfer via automated SFTP)

Då det är en automatiserad överföring som ska köras utan mänsklig inblandning måste en nyckel utan lösenord eller annat skydd av själva nyckeln användas. För att minska risken med detta bör det vara en separat nyckel för endast detta ändamål, och den tillgång som ges bör begränsas i möjligaste mån. Exempel på begränsningar är från vilka IP-adresser den kan användas, vilka kataloger den ger tillgång till och att ingen interaktiv användning där kommandon körs tillåts.

Implementering av UC17 (Authenticated data transfer via interactive SFTP)

Går att implementera på samma sätt som UC01.

Implementering av UC21 (User access to SUPR for changing user attributes not related to authentication)

Implementering av UC22 (User access to SUPR for changing user attributes related to authentication)

Implementering av UC23 (Update of project information in SUPR)

Implementering av UC24 (Direct access to SENSE data storage via sftp (similar to UC14)

Går att implementera med standard ssh tillsammans med Control master. Då krävs bara en tvåfaktor för en session som hålls igång. OM man vill inskränka möjligheten att sessioner hålls igång för länge kan man implementera tidsbaserade tvingande sessions-avbrott.

Implementering av UC25 (SENS web based ThinLinc)

Den normala thinlincklienten (via ssh) är inte användbar om man vill hålla miljön ordentligt "instängd". Så webbklienten med tvåfaktor (TOTP) är den rimliga lösningen.

Implementering av UC26 (Sensitive data query/aggregation via web-based service)