Risk Analysis of TSD - Tjenester for Sensitive Data

Version 5.0
2021-11-11

Espen Grøndahl, Leon du Toit, Gard Thomassen, Haneef Awan, Dagfinn Bergsager
USIT
UiO

1. Introduction

TSD - Services for Sensitive Data (Tjenester for Sensitive Data in Norwegian) was developed by USIT at the University in the period 2011-2014. It is based on a pilot TSD 1.0 which was made 2008-2011. The system offers a secure environment that meets legal requirements regarding privacy and protection of sensitive data. The service is meant to primarily host sensitive research data, but is also used for other data that require additional protection or sensitive data used for other purposes than research.

Since TSD was launched in May 2014 the service has been in continuous development to meet the evolving needs of the targeted user communities.

2. Background

TSD is in full production from May 2014. This risk assessment acts as the UiO TSD risk assessment, and can be used as the basis for the risk assessment every non-UiO institution must perform before using the TSD system.

3. Limitations

This risk assessment addresses TSD and its major components, as well as surrounding infrastructure that might have direct or indirect impact on the solution. The document does not provide details of the single components, unless it is important to describe the overall risk picture. A brief technical description of each component is provided in the attached Whitepaper.

This document is built so that it first presents the overall risk assessment of the basic TSD product, that comes for all TSD projects. Then each appendix documents the extra risk factors associated with the additional component available by “opt-in” for each TSD project administrator.

The basic solution include:

- Login
- Import/export of data
- User administration system
- Self service system
- Virtual computers
- Storage

- Backup
- Nettskjema, widely used opt-in component
- Colossus supercomputer (HPC), widely used opt-in component

The opt-in components include:
- Print functionality
- Consent portal
- Publication portal
- Norwegian EGA
- Use of Elixir AAI
- Use of two-way copy and paste (only for clinical use)

This document does not cover the risks associated with special solutions implemented to address the needs of single projects, but is based on the general usage patterns. Projects that use the system must supplement with their own judgments which are specific to their projects. Although TSD tries to avoid special solutions for single projects, this may still occur. For these cases extra risk factors are not described here, as they only concern a specific project. The general risk assessment for the basic project package has taken into account all risks that are applied to the basic package by the actual existence of the add-ons, and the general rule of thumb is that no add-on is introduced in such a manner that it jeopardizes or changes the basic-component risk assessment.

4. Short Description of the solution

Further descriptions of the solution can be found in the TSD whitepaper and on the service web-pages. Some details cannot be disclosed and are provided only on request.

TSD has a physically and logically closed network with strong access control, and operations are performed using automation tools where possible. The infrastructure consists of virtualized linux and windows environments dedicated to each project, running on a dedicated project VLAN, connected to the central storage facilities. A number of dedicated VLANs are occupied by administrative virtual machines. A cluster of three Active Directory (AD) instances, one physical and two virtual, running on one of the administrative VLANs, is used for user and group administration, and the distribution of host policies. The project machines can only communicate with the resources on their own project VLAN and have no connection with other VLANs or outside networks. Projects do have some openings to the administrative VLANs to enable usage of shared network services and resources. The only network traffic allowed between the administrative machines and the “world” is for the actual user access, logs and license info. VLAN separation will be changed into microsegmentation at a later stage due to scaling needs.

Users access TSD by establishing encrypted communication with dedicated gateway machines. Any login requires Multi Factor Authentication (MFA) using username, password

and a one-time code. Login is either by using https to the web-based self service portals or using the VMware Horizon View framework to access the actual computers. A person having access to two or more different projects will have a dedicated username per project in order to avoid leakage of data between projects. Functionalities that can enable uncontrolled transfer of data to “the outer world” like cut-and-paste functionalities, mapping of local drives, USB forwarding etc, are disabled.

All network traffic into and out of TSD is controlled by the firewall. This includes login and the API-based data transfer.

The TSD HTTPS API allows data import and export. It supports three authentication methods: basic authentication, TSD two factor authentication and BankID authentication. All API clients are verified by TSD staff and are given access to specific projects on a per-need basis. Basic authentication only allows data import and is only allowed from API clients running on machines with specific IP addresses. The API is integrated with TSD’s internal Identity Provider and authentication and authorization system. MFA authentication using TSD credentials and BankID credentials allow for both data import and export, in principle. The same authorization system applies whereby a project administrator has to grant export rights to a project member. In addition to user authentication, applications have to authenticate themselves towards the API, using a revocable time-limited API key. This allows TSD to revoke access to any application at any time, should it be necessary.

Default settings allows only the PI to export data, and all project members to import data, PIs are free to add export privileges to any user(s) in their projects

All the users and administrative personnel must log onto the TSD with MFA authentication. Technicians and sys-admins have a special “operations” user credential so that every operation can be traced and eventually connected to the responsible person.

5. Security requirements

The degree of security is often evaluated along three axes. Confidentiality (C)- ensuring that no one is able to access the data other than those who have legitimate needs. Integrity (I) - ensuring that the data or the code is not manipulated or changed inadvertently. Availability (A)- ensuring that data is available to the right person at the right time.

Data processed in TSD requires a high degree of confidentiality and integrity. Availability is also very important in the sense that the right person shall have access to the right data at all times. As the total number of users is at approx. 7500 this should not be underestimated, but more focus on the “C” and “I” than the “A”.

6. Security assessment

6.1 Method

The risk assessment is performed with the methodology used at the University Center for Information Technology (USIT). It follows the guidelines provided by UNINETT and is based on a collection of risk factors that are assessed based on the probability and consequence.

The risk assessment is carried out by evaluating risk elements that can threaten the confidentiality, availability or integrity of TSD. Risk elements are extracted through assessments made in the design and development phase of the solution and finally evaluated in the TSD Change Advisory Board, a periodic round table with decisional power involving the technical advisors, the administrative leaders and the IT-security officer.

The present risk assessment methodology is an updated version of the previous one, and it is aligned with UNINETT guidelines for risk assessments.

6.2 Members

The assessment is done by:

Espen Grøndahl – IT-Security officer
Gard Thomassen – Service owner
Leon du Toit – Service manager
Morten Werner Forsbring – Technical Coordinator
Haneef Awan - project and financial admin

7. Acceptance Criteria

Security will always be a consideration and a balance between usability and security. The level of usability should be high enough that the possibility to choose a less secure alternative is not wishful. Some risks will therefore be acceptable to get the functionality that users need.

However given the high security profile of the solution very few risks are considered acceptable. There must be low risk associated with inadvertently exposure of sensitive data. If there are security breaches, confidentiality and integrity are prioritized over availability. This will be emphasized in the risk assessment.

7.1 Evaluation Scales

Risk elements are evaluated on a scale from 1-4 in probability, and 1-4 in consequence, where 1 is associated with low probability, or little or no consequence and 4 is very likely and very severe probability or consequence.

In detail, this security assessment adopts the scale suggested by UNINETT to evaluate the probability:

Low (1)	Medium (2)	High (3)	Very High (4)
One time every 10 years or more seldom	One time per year or more seldom	One time per month or more seldom	Weekly

In details, this security assessment adopts the following scale to evaluate the consequences:

	Personal Data	Project Owner	Service Provider (USIT)
Very Severe Consequences (4)	Incident involving unjustifiable lack of security for personal data	- Incident involving loss of data or communication of data to unauthorized parties - Incident leading to irreparable financial losses	
Severe Consequence (3)		- Incident involving loss of data - Incident leading to significant financial losses	Incident leading to substantial and irreparable financial loss or serious loss of reputation / integrity.
Moderate Consequences (2)		Incident involving service with inadequate quality and low availability	Incident leading to substantial financial loss that can be recovered or loss of reputation / integrity due to compromising of infringing information

Little or no Consequences (4)			incident involving loss of trust between the project owner and service provider

Each risk element is associated with a value obtained by multiplying the values associated with the probability and the consequences.

7.2 Risk elements between 1 and 3:

Risk elements between 1 and 4 are considered acceptable either as they are or after the implementation of dedicated measures and/or routines.

7.3 Risk elements between 4 and 8:

Risk elements between 5 and 8 must be evaluated carefully. They might be considered as acceptable for a short period of time, while the necessary mitigating safeguards are being planned.

7.4 Risk elements between 9 and 16:

Risk elements between 9 and 16 are not acceptable. These risks require service interruption and/or must be compensated with manual control and strong routines until the risk is reduced.

8. Description and assessment of the risk elements

Risk elements are listed in the attached excel file. They are numbered with serial number. Here is a description and assessment of elements requiring mitigating measures or special attention (category [4-8]).

Risk element 5

The risk for data leakage between projects as a consequence of intentional action has inevitably high consequences, but low probability. The projects have a dedicated VLAN and no network traffic between project VLANs is allowed. In addition, the access to the resources on a given project is given through group policies. The implemented safeguards have been tested in a penetration test undergone by TSD in summer 2016 and conducted by a well-known security expert. The test was successfully passed as the attacker did not manage to leak data between projects. The implementation has not changed since the penetration test.

Risk element 7

The risk has high consequences, but low probability. Login to TSD requires MFA, in the rare case in which the username and password are guessed, the one-time code is impossible to guess. To get through this barrier it would require social engineering or other “(spear)phishing”.

Risk element 14

The risk that data is exported through some mechanism other than the APIS, is continuously evaluated by periodic testing, and adequate safeguards are implemented as different scenarios are discovered. There is a risk that an authenticated user having access but not export rights might take photos or video of datasets, but this is obvious for all data access, and must be acceptable.

Risk element 15

The risk of someone hacking the TSD API, thereby gaining access to data stored in memory, as it is being transferred over the network. This is mitigated by two factors: 1) unless in the project area, or in an administrative VLAN protected by the firewall, no data is kept in memory unless encrypted, and 2) the API architecture is such that the application server which has access to plain-text data, stored in the project, is protected by several security layers and authentication mechanisms with high level of assurance.

Risk element 16

The risk of damages caused by disloyal sys-admins and technical staff members is reduced by increasing the awareness of the technicians with regards to the importance of the service and the consequences of its failure. In addition, every sys-admin uses private dedicated user credentials in TSD and most of the operations are traceable.

Risk element 19

The event of fire or similar unfortunate events might have severe consequences. Even if the backups of the disks are kept in a different physical location, the disruption of the machine room could result in severe damage of the service infrastructure and long outage. Existing

safeguards consist of fire alarms and monitoring of the machine room. In addition, TSD is investigating a mechanism to make off-site replicas of the data repository and infrastructure.

TSD is working on establishing funding for a major resilience installation for all components in TSD that should and could be running at another physical location as an online live backup of the major components.

Risk element 20

The service is divided over two machine rooms each behind four sets of doors, the first (more external) requires a valid card to enter in off-duty hours while open in normal office hours, while the second, third and fourth requires card and code at all times. Both the card and the code are private to the technicians and are issued by the University of Oslo. The system is not locally controlled, but centrally controlled at UiO. To reduce the risk even further inadvertently TSD are working to implement an extra lock to secure the innermost doors. The entire area including machine-rooms are under video surveillance once entering the second door. However, the combined use of the code and the card significantly reduces the probability of a successful unauthorized access. The fourth “doors”, as there are two machine rooms with TSD-machinery, are only accessible to the ones strictly needing access.

9. Final assessment and follow up

Every point on the previous version of the risk assessment has been rectified in such a way that the risk has been reduced to an acceptable level.

10. Conclusion

According to our assessment, the most serious risk elements have been downgraded to acceptable risks. The solution has been designed from the start to have very high security standards and the present assessment has revealed no weaknesses or backdoors into the system.

TSD has also established a system for incident handling. Cases are reported and archived so that we can easily look at trends and prevent future similar problems. Serious discrepancies will be reported to the end users.

A Council of Changes (in Norwegian “Endringsråd”) has been established to periodically evaluate the security aspects of the present set up and of on-going developments. The Council involves the IT-security officer, the service owner, the service manager, the technical coordinator and the security experts. Main customer representatives are invited, but veto lies with the IT-security officer. The Council supervises any system-changes and guarantees that

the security standards are respected even when the changes are so small that they do not invoke a Risk Assessment upgrade.

Attachments

- TSD Whitepaper TSD
- Risk elements TSD

Appendix 1 - Opt-in service “Print”

There is a possibility of opening for printing documents seen in TSD from the local-host. The project owners of projects enabling this service must consider the following:

- The sensitivity level of data that might be printed (un)intentionally
- The chance of print-outs getting lost (i.e. require pull-print etc)
- What policy and information should be developed for printing in the project

The impact of a print-breach will always be dependent on what (datatype and amount) is printed, and must be assessed by each individual PI. The probability of such a breach will depend on the information, practices and policies that are established for printing from localhost in a TSD project. TSD recommends as little usage of the print-option as possible.

Appendix 2 - Opt-in service “Consent”

TSD supports a digital dynamic consent solution for giving a level-4 (BankID) signature on consents regarding access to data. Risks that should be considered in such cases is the risk of a connection between person identifiers (i.e. name / personal identification number) and some information embedded in the consent that gives for instance health information about the individual. An example would be that there is a leak of a person “Ola Nordmann” participating in a mental health study.
The probability of such a data-breach is small, considering that access can only be gained through TSD or Bank-ID MFA, and the strict access control after this initial authorization.
The impact of such breach will depend on the number of consents involved and the impact factor the knowledge of study participation that the consent gives (for instance mental disease study vs child behaviour in kindergarten).

Appendix 3 - Opt-in service “Publication portal”

TSD has an opt-in service called the “Publication portal”, through which entitled users can grant access to insight or actual download of data that reside inside TSD. Sharing or publication of data should only be granted based on the fact that the sharing and publication is formally okay and needed. Data can be shared so that Norwegians may download or access their data only by using BankID or their TSD-ID. For sharing data to non-Norwegian BankID-holders the receiver of the sharing needs to be a granted user in TSD, i.e have to be a registered user that has been granted access as an “individual to whom data might be shared from project “X” by a project “X” administrator.

The possibility of sharing data wrongfully will always be present and TSD strongly encourages users to be careful, and use this service as automated as possible, and to establish clear routines for who is allowed to share what to whom in what situations.
The impact of wrongfully sharing data from TSD will depend on the amount and type of data.

Appendix 4 - Opt-in service “Norwegian EGA”

TSD is the host of the national node of the Federated European Genome/phenome Archive. This is a service that can receive requests to query datasets in TSD to check whether a given genomic constellation exists in the datasets to be queried. There are risk-factors in such a setup, and risk assessment of NFEGA has been established in addition to this appendix. Please contact TSD or Elixir Norway to see the full risk assessment.

Appendix 5 - Opt-in service “ELIXIR AAI”

In TSD there is a possibility of enabling some user-facing services to users not having BankID/minID or TSD-ID authentication and authorization. Other AAI solutions than the TSD

IAM system will as of today not grant access to a login to TSD, but in some cases there is a wish of for instance allowing an Elixir user to upload data into a given TSD project.
The possibility of data leakage caused by using this opt-in service is considered very low, but there is a risk of someone uploading too much data and thus invoking a denial of service in a given project.
The impact of such a risk factor is considered low.

Appendix 6 - Opt-in service “Token based upload service”

TSD has a service where one can generate an upload token that once/multiple times grants access for a given period of time to upload data to a given TSD project.
The possibility of data leakage caused by using this opt-in service is considered very low, but there is a risk of someone uploading too much data and thus invoking a denial of service in a given project.

Appendix 7 - Opt-in service “Two way copy and paste”

For clinical use at the Oslo University Hospital (OUS) TSD has enabled two-way copy and paste between the TSD virtual clients and the local hospital client. This has been secured by establishing a separate Horizon View connection server that only accepts connections from the projects doing clinical genomics, and that has an IP from a range that is predefined as hospital clinical IT systems. Further there has been an assessment of the risk of someone establishing a man-in-the-middle attack by highjacking / faking the IP-address of the clinical network between OUS and and TSD.
The possibility of a data leak is very low as even with a man-in-the-middle attack, the man in the middle needs to have broken the OUS TSD login MFA login procedure, as the traffic is end-end encrypted over https and the only way to see real data is to complete a full authentication and authorization process towards the designated OUS clinical system Horizon View connection server in TSD.
The impact of such a breach will be limited to the impact of the data that could possibly leak if the connection is compromised.