



SUNET

Sunet Drive



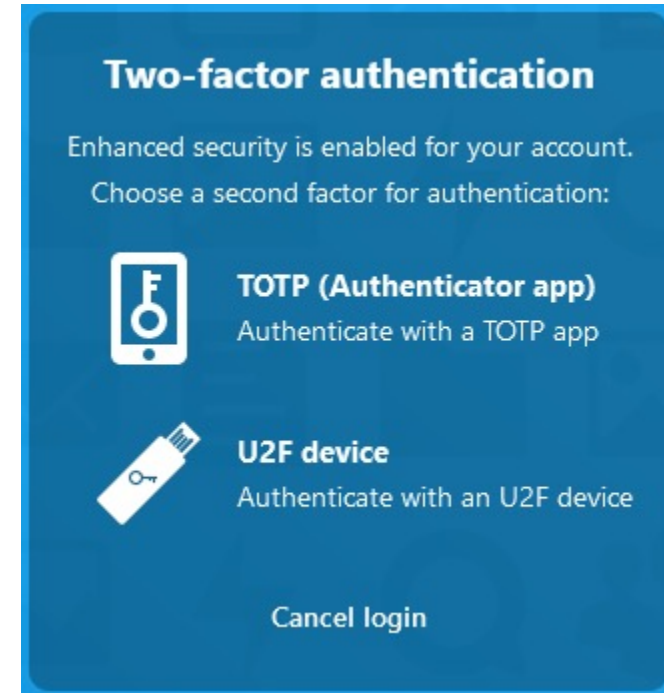
2021-10-27

Sunet Drive – Now and Then

- Multifactor Authentication
 - Why it sounds easy, but is hard to implement
 - Nextcloud step-up authentication
 - IdP-implementation
 - EduID-implementation
- Integration with SND/Doris
 - How would an ideal integration look like?
 - Comparison with other services (Harvard Dataverse, Zenodo)
- Office-integration
- Other features
 - Monitoring
 - Compute-integration
 - edusign
- Storage incident
 - What happened during the incident?
 - Resolution and Outcomes
 - Replication of data
 - Backup, backup, backup

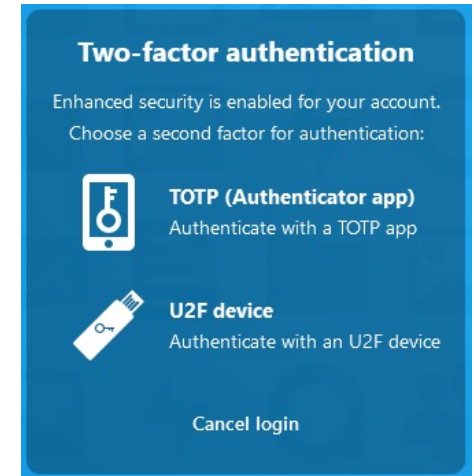
MFA – What is MFA?

- Multifactor Authentication
 - Security technology that requires multiple methods of authentication from independent categories of credentials
 - Examples: SMS, TOTP (Authenticator app), U2F (Yubikey)



MFA – Why is it complex?

- Multiple technologies
 - SMS
 - TOTP (Authenticator app)
 - U2F Device (Yubikey)
 - Other (Mail, proprietary)
- Who implements MFA?
 - Identity provider
 - Service provider
 - Both (?)
- Requires administration
 - Lost or stolen devices
 - Remove/Reenable MFA
- Identity providers
 - Must implement MFA individually
 - Different technologies for different IdPs
- Service providers
 - Can enforce same MFA technology for all
 - Could turn Two-factor into Three-factor if IdP has added factor



Nextcloud and MFA

- Current state:
 - Either SAML/SSO-login or local login (non-SSO)
 - IdP must implement MFA in first case
 - MFA can be added in second case
 - No support for “Step-up Authentication”

Step-up Authentication: Add MFA independently from whether an IdP has MFA or not. Could result in “Three-factor”-Authentication.

Sunet Drive Requirements

- MUST support both MFA and non-MFA logins
- MUST have control over MFA- and non-MFA-areas
- MUST prevent other sharing for MFA-only areas
- SHOULD be easy to administrate

Nextcloud and MFA

- Current state:
 - Either SAML/SSO-login or local login (non-SSO)
 - IdP must implement MFA in first case
 - MFA can be added in second case
 - No support for “Step-up Authentication”

Step-up Authentication: Add MFA to existing login-method. Could result in “Three-factor”-Authentication.

Sunet Drive Requirements


- MUST support both MFA and non-MFA logins
- MUST have control over MFA- and non-MFA-areas
- MUST prevent other sharing for MFA-only areas
- SHOULD be easy to administrate
- SHOULD be easy to use

Conclusion: MFA for Sunet Drive requires custom development for a proper implementation.

Harvard Dataverse/Zenodo

Replication Data for: Temperature and Outgroup Discrimination

Version 1.0



Choi, Donghyun Danny; Poertner, Mathias; Sambanis, Nicholas, 2021. "Replication Data for: Temperature and Outgroup Discrimination", <https://doi.org/10.7910/DVNI/A7EDO3>, Harvard Dataverse, V1

Cite Dataset

Learn about Data Citation Standards.

Access Dataset

Contact Owner

There is a new

January 21, 2021

Dataset Metrics

70 Downloads

Files

Closed Access

Files are not publicly accessible.

Description

High temperatures have been linked to aggression and different forms of conflict in humans. We consider whether exposure to heat waves increases discriminatory behavior toward outgroups. Using data from two large-scale field experiments in Germany, we find a direct causal effect of exposure to heat shocks on discrimination in helping behavior. As temperature rises, German natives faced with a choice to provide help to strangers in every-day interactions help Muslim immigrants less than they do other German natives, while help rates toward natives are unaffected by temperature. This finding suggests that there may be a physiological basis for discriminatory behavior toward outgroups.

Subject

Social Sciences

Notes

data and code to reproduce all tables and figures in the main text and appendices

Files Metadata Terms Versions

Search this dataset...

Filter by

File Type: All Access: All

1 to 7 of 7 Files		Download
	DataAnalysis_Temperature.pdf Adobe PDF - 273.0 KB Published Jan 17, 2021 12 Downloads MD5: af6...25f	
	DataAnalysis_Temperature.R R Syntax - 10.1 KB Published Jan 17, 2021 10 Downloads MD5: b27...0e9	

Example Water Tale

Willis, Craig

Demonstration of how to use Whole Tale to develop custom analysis and visualization for data published externally via DataONE. See https://wholetale.readthedocs.io/en/stable/users_guide/quickstart.html for more information.

Run this Tale on Whole Tale by clicking here.

Preview

5e27210e79e7a662d6024757.zip	
5e27210e79e7a662d6024757	
README.md	457 Bytes
bag-info.txt	241 Bytes
bagit.txt	55 Bytes
data	
LICENSE	170 Bytes
workspace	
.ipynb_checkpoints	
wt_quickstart-checkpoint.ipynb	2.8 kB
apt.txt	4 Bytes
postBuild	28 Bytes
requirements.txt	17 Bytes
wt_quickstart.ipynb	29.4 kB
fetch.txt	114 Bytes
manifest-md5.txt	480 Bytes
manifest-sha256.txt	736 Bytes
metadata	
environment.json	703 Bytes

Files (44.0 kB)

Name

Size

136

views

58

downloads

See more details...

Indexed in

OpenAIRE

Publication date:

January 21, 2020

DOI:

DOI: 10.5281/zenodo.3620753

Keyword(s):

Tale Examples

Related identifiers:

Cites
10.5065/D6862DM8

License (for files):

Creative Commons Attribution 4.0 International

Versions

Version 3

10.5281/zenodo.3620778

Jan 21, 2020

SND/Doris Integration



New data description

Title *

Enter a descriptive, preferably unique, name for the data description, both in Swedish and in English. If no title is available in Swedish, or if it is difficult to translate, the English title can be entered in both fields.

Data accessibility level *

Access to data through SND or other external actor *

Specify where data will be available. You choose whether the data is made available through SND's research data portal, via downloading or by request, or through an external s.. [Show more](#)

- Access to data through SND
- Access to data through an external actor

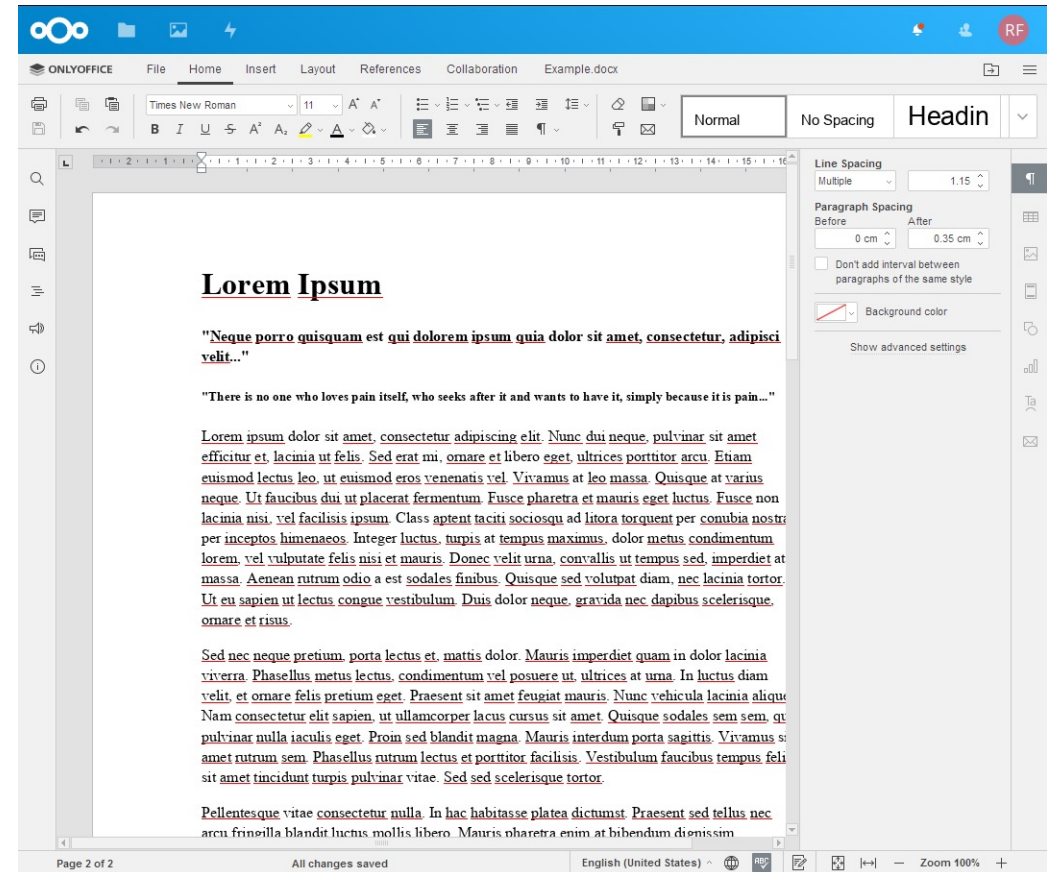
Level of accessibility *

Data that are shared by SND can have two different accessibility levels. Data should, in keeping with the recommendations from the Swedish Research Council, be openly accessible as .. [Show more](#)

- Data are freely accessible
- Access to data is restricted

Office Integration

- OnlyOffice
 - (Relatively) easy to set up for single instance
 - Needs to scale to large instance
 - License cost
- Collabora
 - More complex setup
 - License cost
- MS Office online
 - Medium complexity to set up
 - Unclear licensing situation
 - Microsoft...



Monitoring and Compute Integration

- Monitoring
 - Relatively easy requirements
 - Gets complex when considering individual monitoring setups
- Compute Integration
 - Run computations on top of Sunet Drive storage
 - Unclear requirements
 - Will require a lot of “trial and error”

eduSign

- Simple and straightforward web application
- Requires connection/integration with Nextcloud
- <https://edusign.sunet.se/>

eduSign



Signed in as Richard Freitag (freitag@sunet.se)

Logout



Tap here to choose files to be signed

Storage incident

Summary

On the morning of Thursday, the 17th of June, Sunet Drive reported a lot of gateway timeouts (504) for the object storage in Sto4. The problem was reported at 10:13 CET to Safespring via their support portal. Subsequent investigations led to the conclusion that another customer was unintentionally causing a larger than expected load on the object storage, which as a result had to be taken offline. A detailed technical description of the incident can be found in the document “2021-06-17 sto4 ceph cluster”. As a result, instances of Sunet Drive using Sto4 as their storage backend had to be taken offline and prioritized buckets/customers are at the time of writing waiting for the results of the restore initiative from Safespring.

A main reason for the severity of the incident was the decision to assign backup responsibilities to the customers and end-users of Sunet Drive. This led to the situation where the only copies of certain files resided solely in the affected object storage.

Before the incident, the cluster contained about 9 million objects. When the incident occurred, another customer had uploaded 680 million objects.

Resolution and Outcomes

Resolution

At the time of writing, the restore efforts are still ongoing. All affected end users have been contacted and critical S3-buckets have been prioritized. Due to the timing of the incident before the summer holidays, it is expected that more data/buckets will be prioritized in August. The Sunet Drive pilot environment has been selected as the current replacement for Sunet Drive and is currently being hardened to avoid further incidents. The outcome of this is described in the next section.

Outcomes

The main outcome of the incident is the development and implementation of backup and mirroring of the customer data. This will eventually be done by answering the following question: *“How much data do we potentially lose if a disaster happens in one datacenter?”*. A pragmatic approach based on available skills and technical solutions will be implemented. This means, that certain technologies will not be taken into consideration within the scope of this incident (e.g.: use of CoW filesystems or RAID setups spanning multiple datacenters).

Replication of data

Summary of approach 1: Replication and backup from S3 to S3

Replication from S3 to S3 has the goal to create an identical copy of an S3 bucket in another S3 bucket residing in another datacenter. A third bucket can be used for backup of changed files, essentially resulting in a simple implementation of “copy on write”.

```
rclone sync sto3:bucket1 sto4:bucket1.clone --backup-dir  
sto3:bucket1.backup/Y-m-d_H-M
```

This approach when regularly executed results in a mirrored copy of the data in bucket1 and bucket1.clone, while changed files are being saved in bucket1.backup in a timestamped folder. Essentially, all changed data will be stored and only deleted if actively implemented.

The frequency of the backup needs to be scaled with the amount of data, due to the eventual time to compute the compare operations between the buckets. Depending on the amount of data, the frequency should be minutes to hours for frequently changing data, and hours to days for infrequently changing data.

Backup, backup, backup!

<https://wiki.sunet.se/display/Drive/Sunet+Drive>

<https://wiki.sunet.se/display/Drive/Lagringsincident+sommar+2021>

Contact

- Contact anders@sUNET.se or freitag@sUNET.se
- Test-run using pilot.drive with S3-buckets and more data
- Provision nodes for test and production
- Regular feedback- and training-sessions

Other questions?

Anders Nilsson

Kontaktansvarig för tjänsterna Backup, Lagring och Virtuella servrar

anders@sUNET.se

Richard Freitag

Projektledare

freitag@sUNET.se