2020 IEEE 6th International Conference on Dependability in Sensor, Cloud and Big Data Systems and Application (DependSys) | 978-1-7281-7651-2/20/$31.00 ©2020 IEEE | DOI: 10.1109/DEPENDSYS51298.2020.00026

Investigating Root Causes of Authentication Failures Using a SAML and OIDC Observatory

Jim Basney
NCSA
University of Illinois
jbasney@illinois.edu

Phuong Cao
NCSA
University of Illinois
pcao3@illinois.edu

Terry Fleury
NCSA
University of Illinois
tfleury@illinois.edu

***Abstract*—Authentication is the most critical gatekeeper to the web applications that scientists use to carry out collaborative research. While authentication rarely fails, the impact of failures is huge, and root causes are not well understood. This paper analyzes the root causes of authentication failures from a production authentication system called CILogon, an ideal observatory to monitor authentication issues in a distributed identity federation. CILogon is used by 250+ identity providers and 150+ web applications while acting as a proxy to bridge different single sign-on protocols (OIDC and SAML). Our data on authentication is unique because it is: i) longitudinal (over thirty months), ii) realistic (8,000+ active users), and iii) large-scale (nearly three thousand failures out of 447,428 successful authentications). Our finding is surprising: OIDC has about double the failure rate compared to SAML, which contrasts with our prior belief that SAML is much more complex than OIDC. Our most impactful contribution is a fault tree of error types that quickly finds and mitigates the root cause of authentication errors.**

***Index Terms*—Authentication, distributed systems, error analysis, error handling and recovery.**

I. Introduction

CILogon [1] enables researchers to log on to cyberinfrastructure using their preferred identity providers (IdPs), including more than 3,000 globally distributed campus IdPs in addition to well-known IdPs such as GitHub, Google, and ORCID. The CILogon authentication service is used by researchers worldwide, connecting 8,000+ active users to 150+ research applications using 250+ unique identity providers each month, via the standard Security Assertion Markup Language (SAML) [2] and OpenID Connect (OIDC) [3] web authentication protocols. CILogon also serves as an observatory for Internet-scale use of web authentication for research collaborations. In contrast to common web authentication scenarios that rely on a single IdP, our CILogon service provides operational insights into federated authentication at scale using many IdPs.

This article presents the root causes analyses of authentication failures observed in our production deployment of CILogon. Our analyses are driven by a longitudinal collection of CILogon server logs at the National Center for Supercomputing Applications (NCSA). The CILogon server oversees the authentication attempts of many high-profile, collaborative research projects such as the Laser Interferometer Gravitational-Wave Observatory (LIGO) and the Rubin Observatory Legacy Survey of Space and Time (LSST). Specifically, this article presents a fault-tree characterizing the most common and the most impactful failure scenarios based on an analysis of 2,956 failures over 30 months. These scenarios include SAML faults (18.8%), OIDC faults (36.6%), and background noise, including probes by internet bots (44.2%).

Our key contributions are:

- The first reliability analysis of major authentication protocols (SAML and OIDC) in a large-scale production system.
- A fault-tree that clearly shows the path leading from the root-cause to the observable authentication failures.
- A quantitative analysis of the probabilities as each root-cause branches out in the tree.

II. Background

CILogon began operation in 2010 as a credential translation service [4], accepting SAML authentication from campus identity providers in the US InCommon federation to issue X.509 certificates compliant with International Grid Trust Federation (IGTF) requirements for access to worldwide distributed scientific cyberinfrastructure such as the eXtreme Science and Engineering Discovery Environment (XSEDE) [5] and the Worldwide Large Hadron Collidor Computing Grid (WLCG).

978-1-7281-7651-2/20/$31.00 ©2020 IEEE
DOI 10.1109/DependSys51298.2020.00026

CILogon has since added support for the issuance of OIDC and OAuth tokens [6] for web-based scientific services such as science gateways [7]. Over 8,000 research scientists access CILogon each month to obtain credentials for access to scientific cyberinfrastructure.

While the thousands of campus identity providers supported by CILogon implement the SAML 2.0 standard, they use a variety of SAML software resulting in a wealth of failure cases. CILogon relies on the Shibboleth software to validate incoming SAML assertions, including digital signatures, time constraints, and global namespace constraints. If Shibboleth successfully validates the basic security properties of the SAML assertion, CILogon performs additional validation, such as checking if the SAML assertion contains required user attributes (e.g., identifier, name, and email address) and meets the policy requirements of the service that the user wants to access (e.g., use of multi-factor authentication and conformance to security incident response policy standards). When errors occur, CILogon provides a self-service form allowing the researcher to report the error (with associated technical details) to their identity provider operators using the contact address provided in the federation SAML metadata.

Similarly, while the hundreds of scientific applications supported by CILogon implement the OAuth and OIDC standards, they use a variety of software resulting in another set of failure cases. CILogon has an internal implementation of OAuth and OIDC, which performs the required checks from these standards, including registration of allowed client_id and redirect_uri values, verification of client_secret values, time constraints, and consent for release of user attributes controlled via scope values. When errors occur, CILogon identifies the error cause in a standard OAuth/OIDC error response to the client application.

III. Motivating Example

To enable the exchange of SAML assertions, Identity Providers (IdPs) and Service Providers (SPs) typically rely on third-party software packages to compose the SAML and handle validation of signed and encrypted assertions. One commonly used implementation is available from the Shibboleth Consortium (https://www.shibboleth.net). CILogon uses the Shibboleth SP software to handle the consumption of SAML and the included Apache HTTP Server module (mod_shib) to export attributes to the web server environment for use by web applications. Many Identity Providers use the Shibboleth IdP software to construct SAML assertions containing user attributes consumed by CILogon. CILogon can then assert these user attributes to OpenID Connect (OIDC) Relying Parties (RPs) using CILogon's OIDC Provider (OP).

CILogon identifies users by an eduPersonPrincipalName (ePPN) or a "targeted id"/"persistent id" referenced internally by CILogon as the eduPersonTargetedID (ePTID), following the eduPerson standard (https://refeds.org/eduperson). CILogon expects this identifier to be unique, persistent, and not re-assignable according to the eduPerson standard. If an IdP asserts a different ePTID for a user previously logged on to CILogon, that authentication is flagged as an error (i.e., a possible identity re-assignment). In almost all instances, this error is due to the Identity Provider upgrading the Shibboleth IdP software without properly transferring the secret and hashing algorithm from the previous Shibboleth IdP installation.

Several CILogon OAuth clients (e.g., Globus [8]) also depend on the properties of ePTID and will raise an error when the ePTID for a given user changes. So when an Identity Provider misconfigures a software update, CILogon must contact IdP administrators and assist with proper configuration to assert the original ePTIDs. On a few occasions, it has not been possible to reconfigure the software to restore the original ePTIDs. In these cases, CILogon has had to i) manually update the local database to accommodate the new ePTIDs and ii) contact system administrators of the CILogon OAuth clients and inform them to update their local databases.

The documentation for Shibboleth IdP is considerable, and it can be a daunting task to find the necessary information for seamless upgrading. After feedback was given to the Shibboleth Consortium about the problems Identity Providers experienced upgrading from v2 to v3 of Shibboleth IdP, additional documentation was provided to assist upgrading from v3 to v4. CILogon also developed improved email templates for contacting IdP administrators and CILogon OAuth clients when this particular issue arose. While it is likely neither of these improvements will eliminate this error condition, CILogon is better prepared to handle it if it arises in the future.

IV. Fault tree model

Having described one error case in-depth in the prior section, in this section, we present the cause of each type of fault and develop our fault tree model. All faults result in a common root failure, whereby a user cannot be authenticated to use a

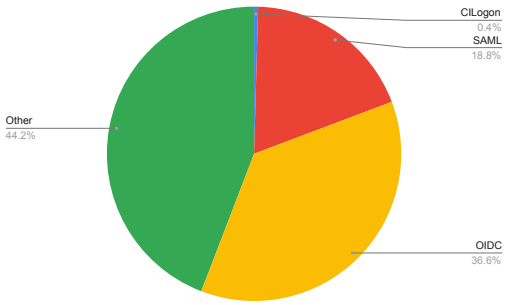

Fig. 1. We attribute error instances to CILogon internal faults (0.4%), SAML IdPs (18.8%), OIDC RPs (36.6%), and Other (44.2%).

cyberinfrastructure service that leverages CILogon for federated access.

A. Dataset, method, and overview of errors

Dataset. We performed our analysis on a 30 month (January 2018 through June 2020) dataset of error log messages by CILogon servers and notification emails from users. The dataset contained 1,646 OIDC error logs, 1,136 SAML error logs, and 173 user emails, for a total of 2,955 faults versus 447,428 successful authentications over the same period. The anonymized dataset is published at https://github.com/cilogon/DependSys2020.

Method. We assigned each error with a unique identifier as the leaves in a fault tree. Then, we grouped the errors into a parent category as the branches in the fault tree.

We performed error analyses only on isolated and mutually exclusive errors. We do not observe any *correlated* errors in which multiple types of errors occurred together in burst, due to CILogon's policy of aborting on the first fault. In cases where an error occurred multiple times in a row from the same IP address (i.e., the user retried a few times), we count the error only once.

Overview of errors. Figure 1 shows the distribution of four main error categories. The majority of errors are due to OIDC configuration (36.6%) and other causes (44.2%). The minority of errors are due to SAML IdPs (18.8%) and CILogon itself (0.4%). In the next section, we identify the specific error cases that make up each of these categories.

Summary of findings. We have the following key findings from our analysis:

1. While OIDC is promoted as a "simple identity layer" alternative to SAML, to our surprise, OIDC accounted for significantly more faults than SAML.

2. As expected, fault rates increased as the number of active users, IdPs, and scientific applications connected to CILogon increased.

3. Faults typically occurred during initial SAML or OIDC endpoint configuration or on subsequent configuration changes.

4. Fault rates tend to track seasonal usage patterns (e.g., dips during the December holidays).

5. Improvements in CILogon OIDC error reporting in February 2020 coincided with a significant drop in OIDC error rates.

6. Large-scale SAML and OIDC authentication can reliably serve thousands of users per month, using hundreds of identity providers and accessing hundreds of applications, with relatively low error rates (under 200 faults per month).

We discuss these findings in the following detailed analyses.

B. Detailed analyses of errors

Figure 2 shows our fault tree consisting of: i) the common failure at the root of the tree (User Unable to Log On), ii) the four main error categories (SAML, OIDC, CILogon, and Other), and iii) the underlying errors at the leaves that cause the failures.

Next, we describe the error cases in the four main branches of the tree.

1) CILogon errors (0.4%): We attribute a small percentage of errors to CILogon software faults, of which there were two causes during the 30 month period we studied:

Database Error: This error occurred once, when the CILogon web application's JDBC connection to the database failed, causing a short outage (with multiple errors in the logs) before CILogon operators initiated a manual fail-over to a secondary production server.

Shibboleth Daemon Error: On four occasions, the Shibboleth software experienced an internal communication failure between the Apache module (mod_shib) and the back-end daemon (shibd), requiring a manual fail-over and restart by the CILogon operators.

Over the 10 years operational lifetime of CILogon, we have certainly experienced more faults than these, including misconfigurations of the CILogon software, implementation bugs, and network outages. These faults have decreased over the years thanks to the maturing of our systems engineering processes and maturing of the CILogon software itself. Over the 30 month period studied, we performed nine scheduled CILogon service updates without downtime.

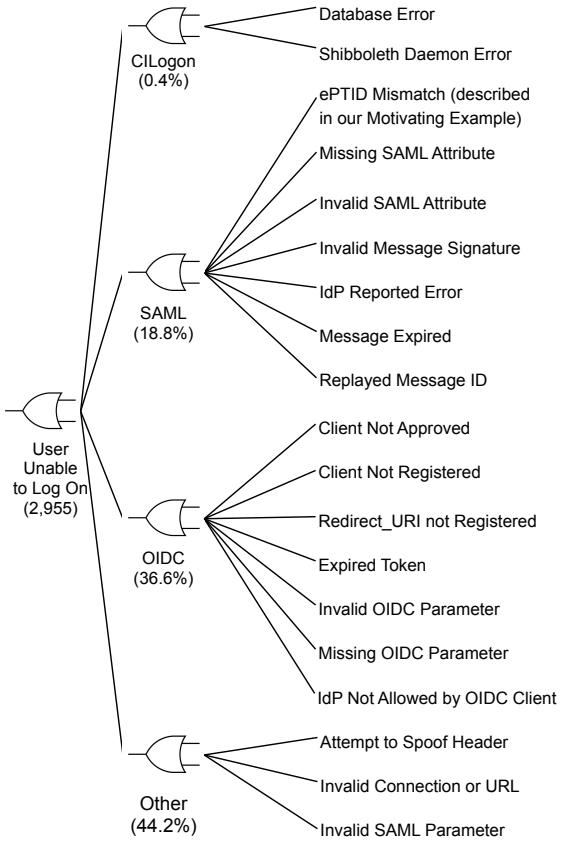

Fig. 2. Our fault tree illustrating: i) the common resulting failure at the root of the tree (User Unable to Log On), ii) the four main error categories (SAML, OIDC, CILogon, and Other), and iii) the underlying errors at the leaves that cause the failures.

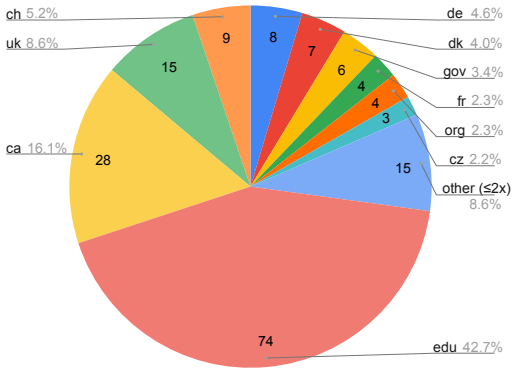

Fig. 3. The Top Level Domains (TLDs) of Identity Providers which did not provide all required attributes, as reported by end users. 11 TLDs (.at, .be, .br, .eu, .fi, .it, .kr, .lb, .no, .pt, and .se) with two or fewer user reports are represented as "other" in the chart.

2) SAML IdP errors (18.8%): One of the surprising (to us) results of our analysis is the relatively small percentage of errors attributed to SAML IdPs, given the complexity of SAML and the differences in federation practices across countries. While we have had growing pains over the years working with a larger number of identity providers, we have benefited over the past 30 months from the maturing of the SAML federations in higher education, as well as improvements to our error handling (see Section VI).

We attributed the SAML IdP errors that we experienced over this period to the following causes:

ePTID Mismatch: This error was described previously in the motivating example in Section III.

Missing SAML Attribute: This error breaks down into 3 cases: 39% of the time, the SAML IdP is not willing to provide required user attributes (unique identifier, name, email address) for privacy reasons; 28% of the time, the IdP attributes in SAML metadata do not meet CILogon policy requirements (e.g., compliance with the federation security incident response policy); and 33% of the time, both required user attributes and IdP attributes are missing. Figure 3 gives a breakdown of "Missing SAML Attribute" errors by Top Level Domain of the IdP, illustrating that, even though 71% of CILogon users come from US .edu IdPs, a significant percentage of attribute issues come from non-US IdPs.

Invalid SAML Attribute: This error occurs when the attribute statements in the SAML assertion do not pass the Shibboleth software verification checks according to the SAML standard. Examples include "Attribute must have Name" and "AttributeStatement must have at least one Attribute". Typically these errors occur when an IdP operator is using CILogon to test a new SAML IdP configuration.

Invalid Message Signature: The typical reason for the Shibboleth software to fail to verify the digital signature of the incoming SAML assertion is a mismatch between the private signing key used by the SAML IdP and the public verification key registered with the SAML federation and used by CILogon. IdPs may update their signing key during a software upgrade or because the existing key is too old, weak, or compromised. To avoid this error, the IdP operator must update the key very carefully, with an overlap period, because there is a lag in time between when the key is registered with the federation and when it is published by the federation and downloaded by CILogon.

IdP Reported Error: This error occurred when

the user's authentication at the SAML IdP was unsuccessful, either due to invalid user credentials or a local fault condition at the IdP.

Message Expired: This error occurs when the SAML IdP's clock is out-of-sync, or there is a 5+ minute delay in delivering the SAML assertion from the IdP to CILogon. According to the SAML specification, CILogon must not accept expired messages for security reasons, though in our experience, these errors are due to misconfiguration or innocent user behavior rather than attacks.

Replayed Message ID: This error occurs when CILogon receives a SAML assertion more than once, which is not allowed by the SAML specification for security reasons. In practice, these errors are due to user browser page reloads rather than attacks.

3) OIDC RP errors (36.6%): We attributed a greater proportion of errors during the 30 month period to OIDC errors. It is our observation that the simplicity of the OIDC specification at-first-glance encourages application developers to implement the OIDC specification themselves rather than using a certified implementation (https://openid.net/certification/). As described below, there are a small number of common error cases that frequently cause trouble when application developers initially connect their OIDC RPs to CILogon.

Client Not Approved and **Client Not Registered**: Client/RP registration is a core concept in the OIDC and OAuth standards. Registering and validating client parameters protects end users against abuse by malicious clients. In particular, CILogon displays the registered client parameters on the consent screen, when the end-user decides whether it is safe to log on to the client and provide their personal information. CILogon provides a web form, and API [9] for client registration. These errors occur when the RP operator has not yet registered their client, has misconfigured their client with an incorrect client_id, or their client has not yet been approved by CILogon operators.

Redirect_URI not registered: The OAuth and OIDC standards require clients to only use pre-registered redirect_uri parameters, to protect against information leakage to malicious sites and open redirector attacks [10]. Invalid redirect_uri values account for fully 49.7% of the OIDC RP errors seen by CILogon, due to RP operators changing the URLs where their applications are deployed. As discussed in Section VI, we improved the error message for this condition in February 2020, resulting in a significant reduction in these errors.

Expired Token: Like SAML, OIDC also makes lifetime restrictions on tokens for security reasons. By default, CILogon issues tokens valid for 15 minutes, so this error typically occurs if there is a long interruption during the user's login flow.

Invalid or Missing OIDC Parameter: These errors occur when the RP omits required OIDC parameters or uses parameter values that are not supported by CILogon's OIDC implementation. This error typically occurs when the RP developer is not using an OIDC certified implementation.

IdP Not Allowed by OIDC Client: This error occurred only once in our 36-month traces, but we note it here for the sake of completeness. CILogon allows RP operators to specify a list of IdPs that are allowed to be used with their application. For example, the application may only support members from specific universities, or the application may not allow Google IDs for privacy/security reasons. Users can only select an IdP from the allowed list at login time, but if the RP operator later removes an IdP from the allowed list, users may still have that IdP selected in their browser session, resulting in this (rare) error case.

4) Other errors (44.2%): **Attempt to spoof header**: This error occurred only five times in our logs, and we believe it was caused by web security scanners probing our site.

Invalid connection or URL: While our site's robots.txt policy disallows all scanning of our site, we still see a constant stream of requests from search bots and other robots probing our site. Invalid URL errors occur when bots scan our site for URLs associated with known web vulnerabilities. Invalid connection errors occur when bots attempt a GET or HEAD request to a SAML endpoint that only accepts POST requests. These errors occur because bots collect our SAML endpoint URLs from published InCommon federation metadata, then scan those URLs.

Invalid SAML Parameter: Similar to "invalid connection" errors, these errors occur when a bot does not provide required parameters (SAMLRequest, SAMLResponse, TARGET, etc.) to our SAML endpoints. While it is possible for these errors to occur due to the misconfiguration of a SAML IdP or due to malicious activity, all the instances we have analyzed indicate SAML-unaware scanning by bots.

V. Error rates over time

So far, in this article, we have examined errors in aggregate for the 30 month trace period. In this section, we look at error rates over time. Figure 4 shows that the number of errors per month is a fairly bursty metric, but the error rate is generally

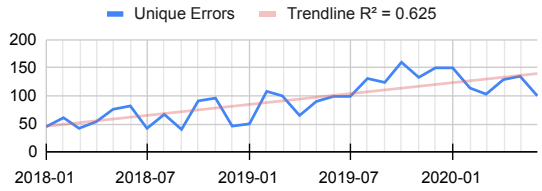

Fig. 4. During the period January 2018 to June 2020, there were an average of 92.7 unique errors with a standard deviation of 35.6. Note there are seasonal dips in errors December-January of each year.

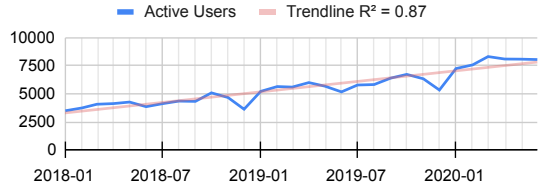

Fig. 5. The number of active users per month increased steadily from 3,488 users in January 2018 to 8,059 users in June 2020.

increasing over time. In Figures 5, 6, and 7, we see that the number of active users, the number of IdPs used, and the number of OIDC clients used are each increasing at a faster rate per month than the error rates. We can also note summer and winter seasonal dips in activity corresponding to the academic calendar.

We also note a significant drop in error counts starting in February 2020. We believe this is likely

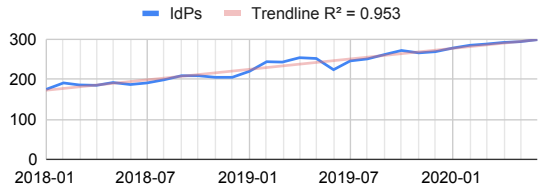

Fig. 6. The number of IdPs per month increased steadily from 175 IdPs in January 2018 to 299 IdPs in June 2020.

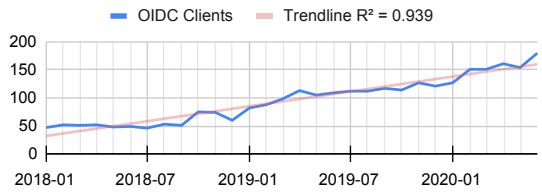

Fig. 7. The number of OIDC clients per month increased steadily from 47 OIDC clients in January 2018 to 180 OIDC clients in June 2020.

due to the change in error handling described in Section VI for the redirect_uri error described in Section IV. Since this specific error occurs at a high rate (49.7% of total OIDC errors and 18.2% of total errors overall), the mitigation for this one case can have a significant overall impact, enabling OIDC RP administrators to resolve this problem more quickly without generating multiple error logs on our side.

Most importantly, we had found no statistical significance correlation, i.e., correlation coefficient = 0.396, among the errors and the active user. This statistic means CILogon handles the growth of users well and can scale to support larger scientific missions.

VI. Mitigation of Errors

When an error occurs during the user's logon, it is often possible to provide the user information, which helps them diagnose and fix the issue on their own. The type of user depends on the error in question.

OIDC client errors typically occur during the initial relying party installation and configuration phase. In this case, the user seeing errors is the administrator performing the configuration of the OIDC client. When CILogon first added OIDC Provider support, only the basic errors were caught and reported, such as "missing client_id". It was incorrectly assumed that OIDC administrators would have a sufficient understanding of the OIDC specification to configure the client correctly. Thus other errors were reported as a generic error message "failed to initialize OIDC flow". Over time, CILogon categorized these general errors into more specific error messages such as "the redirect_uri does not match a registered callback URI." This error description enables OIDC client administrators to more easily debug configuration errors. CILogon will continue to evolve and report more specific errors. A future update will target less common errors, such as "invalid characters (+) in scopes."

Once a client has been configured correctly, errors are instead seen by the end-user, for example, when an IdP fails to release all required user attributes. In this case, the user's authentication at the IdP may succeed, but CILogon would report an error. CILogon has always reported the missing attributes to the user and provided an email address of an IdP administrator to contact about the issue. However, this email address was presented as a `mailto:` protocol hyperlink which would populate a new email message with information about the issue for the IdP admin, as well as a link to

instructions on a CILogon FAQ on how to configure the IdP to support CILogon. Text on the web page also said that users could contact CILogon at the email address in the page footer. Many users preferred to compose an email to CILogon from scratch rather than click on the hyperlink for their IdP admin. This would require CILogon admins to contact the IdP admin on behalf of the user. In February 2020, the CILogon user interface underwent a major overhaul using Bootstrap CSS to simplify the look and feel of the website. As part of this update, the hyperlink for the IdP admin was changed to an HTML button, but the functionality was unchanged. With this simple UI change, we expect more users to contact the IdP admin directly (with a CC to CILogon) instead of composing an email to CILogon.

For those rare occasions when an error is due to a problem with a component of CILogon, Nagios and Monit alerts that have been added to quickly notify CILogon administrators via email and SMS. The current production configuration requires manually specifying which server is “live" for the service. Future plans include moving components of CILogon to cloud providers. This move would also allow for the detection of problems and automated fail-over to another production server.

VII. Related Work

This section compares our work to related work on monitoring authentication services.

Internet-scale observatory. The challenge of operating Internet-scale network services is to: i) measure their availability and ii) quickly respond to outages. The problem is hard because of the globally distributed nature of a network’s components. To address this problem, past work has developed distributed monitors for almost all major network protocols, including Domain Name Resolution (DNS) [11], Secure Socket Layer (SSL) observatory, Secure Shell (SSH) honeypot [12], IoT [13], and censorship [14]. Despite that, no observatory exists for monitoring SAML or OAuth/OIDC.

The vantage point of our observatory is the production CILogon servers, stationed at the National Center for Supercomputing Applications (NCSA). CILogon plays a central role in safeguarding access to scientific resources by coordinating with geographically-distributed IdPs and OIDC clients. Therefore, CILogon provides a centralized database of rich operational logs that reveal authentication failures that involve multiple IdPs and OIDC clients in real-time.

Authentication bugs in OAuth implementations. The problem of correctly implementing authentication logic is hard. Prior work [15] has shown that authentication bugs are prevalent in almost all implementations of the OAuth protocol. These bugs have been actively exploited by attackers in order to take over or to execute actions (e.g., transfer money or buy goods) on legitimate accounts. However, no prior work has analyzed isolated or correlated failures of authentication at run-time.

Our paper focuses on run-time authentication failures that were reported by legitimated users rather than being actively exploited by attackers. In contrast to static analyses of implementation code, our work relies on notifications from emails and operational logs to perform post-mortem analysis and characterization of these failures.

Failure localization and recovery in distributed systems. Cascading failures in distributed systems are inevitable. However, the problem of localizing the root-cause of such failures is hard, mainly due to the inherent large-scale and global deployment of such distributed services. Existing work in other domains such as mobile applications, Windows OS, and cloud systems attempted to characterize failure and used hypothesis testing to find the root cause [16]. However, prior work does not include the authentication failure domain.

Our paper analyzes misconfiguration errors and operational failures rather than bugs that can be exploited intentionally. The causes of these failures have been depicted in the fault tree and can be used to develop more detailed error reporting. Specifically, CILogon sanitizes for malformed input, e.g., invalid message signatures, expired tokens, or invalid redirect URIs, to fail early to prevent errors from propagating in its services.

System misconfiguration. System misconfiguration is the main source of critical failures in production, primarily because of too many knobs for administrators to fine-tune [17]. Our paper confirmed the same configuration challenges with SAML IdP and OIDC RP operators. Some of the most common misconfiguration errors are missing SAML attributes, ePTID mismatch, and invalid client_id or redirect_uri parameters. These issues have been actively addressed by educating InCommon members (via email list and meetings) and improving OIDC error messages. Our refined error reporting mechanisms helped users to self-diagnose and fix these misconfiguration issues. As a result, we saw a decreased error rate starting in February 2020.

VIII. Conclusions and Future Work

In conclusion, we have used CILogon as an observatory to look into faults that occur in a global ecosystem of SAML IdPs and OIDC RPs serving academic research applications. Contrary to the claimed simplicity of OIDC over SAML, we found almost twice as many instances of OIDC errors as SAML errors. We also found that improved error messaging can significantly reduce error rates, as problems are diagnosed and resolved more promptly. Even when including errors from automated probes by internet bots, we show that error rates are relatively low (0.7%) for a service that federates 250+ identity providers with 150+ research applications. While in this article we used CILogon as an observatory to study errors in a federated authentication system, we envision future work using CILogon to study performance (latency, scalability, etc.) and security (cryptographic algorithm choices, software patching practices, incident response policies, etc.) across SAML IdPs and OIDC RPs.

Acknowledgments

This material is based upon work supported by the National Science Foundation under grant numbers 1535070, 1547249, and 1548562. The authors thank the anonymous reviewers for their input which helped us improve this article.

References

[1] Jim Basney, Heather Flanagan, Terry Fleury, Jeff Gaynor, Scott Koranda, and Benn Oshrin. CILogon: Enabling Federated Identity and Access Management for Scientific Collaborations. In *International Symposium on Grids and Clouds (ISGC)*, 2019. https://doi.org/10.22323/1.351.0031.

[2] S Cantor, J Kemp, R Philpott, and E Maler. Assertions and protocols for the OASIS Security Assertion Markup Language (SAML) v2.0. Technical Report saml-core-2.0-os, OASIS, 2005. http://docs.oasis-open.org/security/saml/v2.0/.

[3] N Sakimura, J Bradley, M Jones, B de Medeiros, and C Mortimore. OpenID Connect Core 1.0. Technical report, OpenID Foundation, November 2014. http://openid.net/connect/.

[4] Jim Basney, Terry Fleury, and Jeff Gaynor. CILogon: A Federated X.509 Certification Authority for Cyberinfrastructure Logon. *Concurrency and Computation: Practice and Experience*, 26(13), 2014. https://doi.org/10.1002/cpe.3265.

[5] J Towns, T Cockerill, M Dahan, I Foster, K Gaither, A Grimshaw, V Hazlewood, S Lathrop, D Lifka, G Peterson, R Roskies, J Scott, and N Wilkins-Diehr. XSEDE: Accelerating scientific discovery. *Computing in Science Engineering*, 16(5):62–74, Sept 2014. https://doi.org/10.1109/MCSE.2014.80.

[6] D. Hardt. The OAuth 2.0 Authorization Framework. RFC 6749, IETF, October 2012. http://https://doi.org/10.17487/RFC6749.

[7] Jim Basney, Rion Dooley, Jeff Gaynor, Suresh Marru, and Marlon Pierce. Distributed web security for science gateways. In *Proceedings of the 2011 ACM Workshop on Gateway Computing Environments*, GCE '11, page 13–20, New York, NY, USA, 2011. Association for Computing Machinery. https://doi.org/10.1145/2110486.2110489.

[8] S Tuecke, R Ananthakrishnan, K Chard, M Lidman, B McCollam, S Rosen, and I Foster. Globus Auth: A research identity and access management platform. In *IEEE International Conference on e-Science (e-Science)*, pages 203–212, Oct 2016. https://doi.org/10.1109/eScience.2016.7870901.

[9] J. Richer, M. Jones, J. Bradley, M. Machulak, and P. Hunt. Oauth 2.0 dynamic client registration protocol. RFC 7591, IETF, July 2015. https://doi.org/10.17487/RFC7591.

[10] T. Lodderstedt (Ed.), M. McGloin, and P. Hunt. Oauth 2.0 threat model and security considerations. RFC 6819, IETF, January 2013. https://doi.org/10.17487/RFC6819.

[11] Pawel Foremski, Oliver Gasser, and Giovane CM Moura. DNS Observatory: The big picture of the DNS. In *Proceedings of the Internet Measurement Conference*, pages 87–100, 2019. https://doi.org/10.1145/3355369.3355566.

[12] Phuong M Cao, Yuming Wu, Subho S Banerjee, Justin Azoff, Alex Withers, Zbigniew T Kalbarczyk, and Ravishankar K Iyer. CAUDIT: Continuous Auditing of SSH Servers to Mitigate Brute-Force Attacks. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*, pages 667–682, 2019. https://www.usenix.org/system/files/nsdi19-cao.pdf.

[13] Harm Griffioen and Christian Doerr. Examining mirai's battle over the internet of things. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 743–756, 2020.

[14] Ram Sundara Raman, Prerana Shenoy, Katharina Kohls, and Roya Ensafi. Censored planet: An internet-wide, longitudinal censorship observatory. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, pages 49–66, 2020.

[15] Eric Y Chen, Yutong Pei, Shuo Chen, Yuan Tian, Robert Kotcher, and Patrick Tague. OAuth demystified for mobile application developers. In *Proceedings of the 2014 ACM SIGSAC conference on computer and communications security*, pages 892–903, 2014. https://doi.org/10.1145/2660267.2660323.

[16] Ramnatthan Alagappan, Aishwarya Ganesan, Jing Liu, Andrea Arpaci-Dusseau, and Remzi Arpaci-Dusseau. Fault-tolerance, fast and slow: exploiting failure asynchrony in distributed systems. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 390–408, 2018. https://www.usenix.org/conference/osdi18/presentation/alagappan.

[17] Tianyin Xu, Long Jin, Xuepeng Fan, Yuanyuan Zhou, Shankar Pasupathy, and Rukma Talwadker. Hey, You Have Given Me Too Many Knobs! Understanding and Dealing with Over-Designed Configuration in System Software. In *Proceedings of the 10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE'15)*, Bergamo, Italy, August 2015. https://doi.org/10.1145/2786805.2786852.