Securing Collaborative Work in Wide-band Display Environments

Krishna Bharadwaj
Electronic Visulaization Laboratory, Department of Computer Science,
University of Illinois at Chicago,
Chicago, USA
kbhara5@uic.edu

Andrew Burks, Andrew Johnson, Lance Long, Luc Renambot, Maxine Brown
Electronic Visulaization Laboratory, Department of Computer Science, University of Illinois at Chicago,
Chicago, USA

Dylan Kobayashi, Mahdi Belcaid, Nurit Kirshenbaum, Roderick Tabalba, Ryan Theriot
Laboratory for Advanced Visualization & Applications, Hawaii Data Science Institute, University of Hawaii at Manoa,
Honolulu, Hawaii, USA

Jason Leigh
Laboratory for Advanced Visualization & Applications, Hawaii Data Science Institute, University of Hawaii at Manoa,
Honolulu, Hawaii, USA
leighj@hawaii.edu

***Abstract*—SAGE2 (Scalable Amplified Group Environment) is the de facto platform to support group work on wide-band display environments. Unlike most web applications, the SAGE environment, due to the nature of its collaborative model, needs a nuanced handling of security aspects. This paper details the security requirements of SAGE2, the Identity and Access Control model that was developed to address those requirements, and the details of the Identity and Access Management system that the SAGE team implemented based on this new model. Further, we present a comparison of this new system with some of the popular collaboration platforms to highlight the uniqueness of SAGE2 integrated with this new Identity and Access Management system.**

Keywords—security, access control, collaborative environment, wide-band displays

I. Introduction

Large high-resolution displays (generally called wide-band displays) are being increasingly used in classrooms for teaching, in offices to replace projectors, by analysts on their desktops for data exploration and visualization, and so on. Wide-band displays offer screen real estate that makes them well suited for working with multi-applications tasks [2] as well as multitasking [1]. Adoption of such displays for group work has thus been a natural consequence where collaborations span multiple sessions spread over several days, weeks, months, or even years, and involves generating and sharing multiple documents and data sets.

SAGE (Scalable Adaptive Graphics Environment) [3] and its successor SAGE2 (Scalable Amplified Group Environment) [4, 5] are today's de facto platforms to support such group work on wide-band displays (Examples shown in Fig. 1), enabling users to collaborate locally and remotely, and to access, display, share and manipulate documents and visualizations [6].

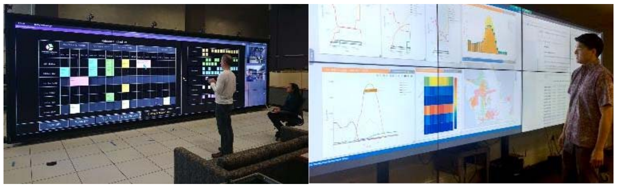

Fig. 1. Examples of SAGE2 wide-band displays

SAGE is an open-source middleware that provides multiple users with a common operating environment to access, display, and share an assortment of data intensive information. The software allows each user to create a pointer on the wide-band display by using their own personal device, or to directly approach the wide-band display and interact through a multi-touch interface. In this manner, multiple users can simultaneously add and interact with content. SAGE displays pixel streams from remote sources by utilizing high-speed networks to render content ranging from high definition images and videos to PDF documents and laptop screens. [4]

For most web applications, security involves encrypting communications, securing resources and providing authorized access to those secure resources. The SAGE environment, due to its open collaborative model (SAGE environment relies on social etiquette for control and manipulation of windows and content on the wide-band display), needs a more nuanced handling of security aspects. Most resources, such as documents, images and so on, in a SAGE environment are user generated. Hence, security and privacy requirements vary with each resource. Also, most of the resources are generated or shared during collaborative sessions, making the security requirements time sensitive as well. Apart from the user generated resources, we also have resources such as storage and computation that require authorized access. Further, the diversity and distributed nature of the SAGE2 user community brings cybersecurity concerns into the mix – whether in the laboratory or the classroom – such as managing who gets to participate in which collaborative sessions and what files they are permitted to share, view and/or download.

SAGE and SAGE2 are funded by National Science Foundation awards ACI-1441963, ACI-1339772.

Since wide-band displays come with a non-trivial cost in terms of money, installation, setup efforts, and take up large physical space, organizations prefer to install a few of these large displays in the premises and make them available to multiple groups across the organization that might need wide-band displays for their group work. This is done in order to maximize the utilization of wide-band displays by interleaving the time spent by those various groups accessing the wide-band displays. With the collaboration of each group spanning multiple sessions, scheduling work on these displays can also be an issue. Another practical concern is that of data compartmentalization. For instance, classes of multiple courses at a university that meet periodically at the same wide-band display (room) for teaching or laboratory work each generate their own content such as student assignments, reading material, and so on. Such content needs to be properly compartmentalized so that the students of one class have access to the content of exactly that class alone. SAGE2 architecture provides a simple way to upload and share content and in its current state does not support such private partitions of data. All these aspects make SAGE a unique security use case for cyberinfrastructure, especially given SAGE's distributed nature and application diversity.

As the first step, the SAGE team sought to understand the security requirements of the SAGE2 environment in detail. Consequently, the team designed an Identity and Access Management (IAM) model appropriate for SAGE's user community. Following this, we implemented an access control mechanism based on the IAM model, that combines layers of management such as calendars and scheduling of events with security aspects such as access control and identity management to address the security requirements of SAGE2. The main contribution of this paper is a novel access control model for supporting group work in wide-band display environments. In this paper we will discuss: (1) User requirements with respect to security concerns in wide-band displays. (2) Some of the main access control models in use, highlighting the reasons that motivated us to come up with a novel access control model. (3) Related work. (4) IAM model that was designed to support the requirements. (5) Identity and Access Management System based on the IAM model and its implementation details. (6) A few added benefits that come with the implementation of IAM model. (7) An evaluation of the system.

II. End User Requirements

The SAGE2 Team began working with Trusted CI, the National Science Foundation's Cybersecurity Center of Excellence, in July 2018 to better understand and address the interactions between humans and security, and to explore how SAGE2 software can best facilitate users taking a larger role in handling security issues unique to their sites. To more fully understand both the security needs of the SAGE2 user base and the impact that IAM integration could have upon the SAGE2 user community, we determined that user feedback was needed. The Teams developed and distributed an online survey in December 2018. Questions focused on how their institutions dealt with identity management, what their concerns were regarding SAGE2 identity and access management, and what were their known SAGE2 user roles and meeting scenarios.

The following summarizes our best understanding of the community's access control requirements. While the original goal was to characterize SAGE2's requirements, we believe the model is broad enough to encapsulate the needs of future cyber-enabled collaboration environments.

In our SAGE2 access model, there are *users*, physical display *walls*, *meetings*, *assets* (documents) and *cyberinfrastructure-services* (CI-services) used to support meetings. *Users*: Users need to be identified, much like a social security number or a passport number identifies an individual. A user's identity determines which walls and meetings they have access to. A user's status differs from meeting to meeting, such as: a wall administrator, a meeting organizer, a meeting participant, a guest, or a spectator (a guest who can only watch a meeting). A user has an affiliation, e.g. University of Hawaii, College of Natural Science, etc., which could determine their access to meetings and meeting resources. *Meetings*: Meetings may occur regularly, or they may have start and end dates and/or times. Meetings have sessions, which can be persistent or cleared after each meeting. Sessions store the complete state and timeline of the meeting, such as documents shown at a meeting and their placement on the wall. Sessions can provide the basis for audits. Meetings have invitees. Attendees may have different access permissions for a meeting depending on their identity; e.g., permission to upload a document, permission to point at something, permission to organize documents on the wall, or permission to download information from the wall. Meetings can be moved to walls in different rooms, hence meeting sessions and their affiliated assets (documents) have to be portable, too. There are various types of meetings, as noted in [7]: brainstorm meetings, data collection, data analysis, knowledge crystallization, knowledge presentation, and classroom lecture. *Assets*: Assets (documents and folders) refer to any form of data used in support of a meeting. Assets may be viewable, writeable, and/or downloadable. The accessibility of an asset can be determined by meeting type, users and/or affiliations, and/or physical walls. For example, a particular document may only be viewed on a specific display wall in a specific room. Assets may or may not persist after a meeting, or they may have a life span; e.g., they could be deleted immediately after a meeting or after a certain amount of time has expired. The life cycle of the asset (creation, views, downloads, etc.) is tracked to provide a reliable provenance history. *CI-services*: CI-services refer to compute clusters, data servers, computational notebooks, video conferencing, and Cloud document services (such as Google Docs), which are accessible from the walls.

Note: Not all users require the full set of access requirements. The users with the greatest restrictions tend to belong to private companies, the military, and some government agencies. Universities and K-12 institutions have comparatively fewer restrictions.

III. Related Work

In this section first we look at some of the main access control models that are in use and discuss their applicability for an environment such as SAGE2, listing the main reasons that motivated us to come up with a novel access control model. Furthermore, we present a few examples of research efforts that attempt application of access control models to collaborative systems. Mandatory Access Control (MAC) [17] defines security levels for both assets as well as users and allows a user to access an asset only if the user has a security level equal to or higher than the asset being accessed. It is a centralized access control model where a security administrator manages the access levels of all assets and users through security policies. Users are not allowed to override these policies. This makes sharing assets a user owns with other users difficult if the intended recipients do not have the same or higher level of access compared to the assets being shared. In an environment like SAGE2 most of the assets are user generated specifically to be shared within a group. Furthermore, the lifespans of meetings, events, assets, and even user accounts fall into a broad spectrum and one time access and ephemeral assets are commonplace. Hence, the MAC model is not well suited for SAGE2. Discretionary Access Control (DAC) [17] in the traditional sense operates by letting users decide who has what level of access over the assets they own/create. While this simple model allows for a lot of freedom in deciding the access for assets, it can lead to long access control lists if the number of users and the number of assets they own are large should the access control be defined at individual user level. More importantly, by making every user responsible for controlling access to the assets they own, the DAC model fails to ensure group level isolation of assets generated by the group. Clearly DAC alone cannot be employed to address the nuances of security requirements in SAGE2. Role based access control (RBAC) restricts access to information based on roles of individual users [9], however RBAC is policy-neutral since the notion of a role is abstract. In the RBAC model, each user in the system can be assigned a role or multiple roles and the definition of a role then determines whether a user in that role has access to a particular asset. We can define as many roles as we need and assign them to users appropriately. While this model seems like a good candidate for SAGE2, there are two main issues that led us to look for a better alternative. First, in an environment that has SAGE2 as the main platform for collaboration, there can be many different types of groups that avail the services of a SAGE2 display, with each group having its own unique group dynamics that call for very group specific roles. It is not practical for a security administrator to conceive all the different roles that every group might need. A solution to this issue could be that the event owner requests the security administrator to create the set of roles needed for an event either at the time of creating an event or as the need arises. However this would induce the overhead of someone(security administrator) having to create/approve the roles thereby inducing a delay in the process because of the injection of a human into the loop. Second, RBAC does not take into account any attributes of resources or of the environment in dispensing access permissions. This means that we would have to look for additional mechanisms to handle scenarios like time sensitive access to assets and so on. Attribute Based Access Control (ABAC) [8] offers a granular approach to access control which makes it very powerful. It allows access control to be configured based on attributes of not just the assets and the users trying to access those assets but also attributes of the environment itself such as date, time, place, and so on. However for a system like SAGE2 which is heavy on real time user interaction and accessing an asset mostly translates to putting it on the large display and interacting with it on the large display, most of the access control decisions need to happen under the hood so as to keep interruptions to a minimum if we are to ensure a smooth user experience. Hence the decisions, users at any level of access have to make, need to be minimized. This prompted us to come up with a model that exploits the granularity offered by the ABAC model and provides a way for users to manage access within their own groups without much interruptions in their workflow or without having to rely much on the security administrators. In fact our proposed model incorporates the strengths of both RBAC as well as DAC in providing the users with a secure system for group work. Within an event, it exposes an RBAC model allowing the event owner to manage access control of the group and at the same it follows the DAC model in essentially making the event owner responsible for deciding what happens to the content that is generated during the course of the event.

We now present a few examples of application of access control models to collaborative systems. Virtual Organization Membership Service (VOMS)[12] is an authorization system that deals with granting users access to resources in a computational grid. By creating a virtual organization which is an abstract entity grouping users, institutions and resources in the same administrative domain, VOMS addresses the issue of granting access to remote resources to users who might not belong to the same site where the resource is located without having to repeat authorizing users for every remote resource. VOMS concerns itself with remote access issue whereas our proposed system is trying to manage access control in a collaborative environment that has both collocated users as well as remote users trying to access commonly visible data on the large display as well as other resources made available by the SAGE2 server. INDIGO-Identity and Access Management (IAM) provides IAM as a service to manage identities, enrollment, group membership, and policies to access distributed resources in a homogenous way[13]. Even though it addresses a few of the security aspects targeted by our proposed system, since INDIGO-IAM offers them as a service, they cannot readily be applied to the realm of SAGE2 without the requisite event management infrastructure that our proposed solution offers. Furthermore, INDIGO-IAM deals with policies related to distributed resources whereas access control in SAGE2 is more nuanced as it deals with moderating user interaction with contents on the large display apart from controlling access to resources in the traditional sense. WiseShare [18] is a social collaborative environment built on ABAC policies for secure sharing of knowledge. Based on the reliability of the content a user produces, the user is tagged, and appropriate privileges are given to the user. The more reliable content, a user produces, the better user's privileges are. The tool itself is limited to sharing of documents between users and lacks any notion of roles or their customization unlike the IAM model that we present. Nasirifard et al. [19] present an annotation-based access control model for collaborative platforms. The model uses the concept of "tagging" to define customized roles thereby providing discretionary access control, however it does not consider any contextual information and relies purely on user created tags. A dynamic access control model is presented by [14] that relies on task role based access. In this model users gain access to resources based on roles assigned to them and these roles can be dynamically changed depending on the task they are working on. Even though this model allows user roles to be changed dynamically similar to the IAM model that we present, unlike our model, it does not support discretionary access control and relies on a user of authority such a security administrator to execute the change of roles.

IV. Identity and Access Management System

The IAM model was developed to address the access control requirements of SAGE2 described in the previous section. The model is designed to capture the interconnections between the different components of the SAGE2 access model such as walls, meetings, assets, and CI-services. The IAM model is based on the notion of events. An event models the way groups meet for work around wide-band displays, incorporating the secure access requirements surrounding the assets and CI-services. An event is a set of time slices (a single session, or a recurring set of sessions with a defined start and end date) on the SAGE2 server. Each session is an uninterrupted run of the SAGE2 server facilitating the collaboration. The basic idea is to let authenticated users schedule access to wide-band displays for their group of users and carry out their collaborative work during the scheduled sessions in a secure way. The decision to model the access control mechanism based on events as opposed to a more open abstraction scheme is mainly to have this new secure model be accessible across the spectrum of SAGE2 user community. SAGE2 users come from different backgrounds such as private corporations, universities, government agencies and so on. Having an open abstraction scheme would force the users to come up with their own notions of meetings and incorporate that into the SAGE2 system through further development. While private corporations and such big facilities might have the resources to accomplish that, many single system users and smaller establishments might not be well equipped to invest in such exercises. A centralized server (shown in Fig. 2) called the Identity and Access Manager (IAManager) manages the various SAGE2 installations within an organization by facilitating event creation on those servers, thereby managing access control. User authentication is carried out using OpenID Connect through CILogon. CILogon is an integrated open-source identity and access management platform for research collaborations [10]. It supports over 3000 identity providers. CILogon provides a standards-compliant OpenID Connect [11] (OAuth 2.0) interface to federated authentication for cyberinfrastructure (CI). IAManager hands off the authentication to CILogon and upon successful authentication, the users are allowed access to the IAManager. Once in, the users can then gain access to all the SAGE2 servers within the organization. That is, a user can create events on any SAGE2 server (i.e., the corresponding wide-band display).

By creating an event a user "blocks" or "reserves" access to SAGE2 server during the scheduled sessions of the event. Through the event, users can interact with the wide-band displays by adding content and applications onto the display. Event owners can add other users to the event thereby making them part of the group interaction with the wide-band display. Different members of the event can be assigned different roles which will determine the different interactions that users can perform on the wide-band display, from simple window manipulation to content creation and manipulation, and even some meta interactions like changing the roles of other users. Event owners can create custom roles for every event. A fixed set of roles fails to capture the possibility of a user whose set of access permissions lie in between that of available roles. Custom roles enable for a more granular access control. Furthermore, event owners can choose role names that are relevant to the context of the event making it easier to keep track of the access permissions of these roles and their assignments to users. For instance, the teaching assistant (TA) of a course could be assigned a role named "TA", whereas a project meeting event in an office could have a role named "Presenter" that could be assigned in each session, to the person who will be presenting in that session. Content generated during an event stays private to that event and is accessible only when the event is in session. We employ Attribute Based Access Control model (ABAC) to realize the events in SAGE2.

The Identity and Access Management System not only provides *users* secure access to *assets* and *CI-services*, but also makes it easy to schedule work on the *walls* and manage *meetings*. The IAManager manages the authentication of users and acts as a gateway to different SAGE2 servers within the organization. A user who logs into the server can potentially (upon authorization) access any SAGE2 instance, or in other words, any wide-band display within the organization. As shown in Fig. 3, for identity management, it employs OpenID connect to offload the authentication to identity providers via CILogon's OpenID connect module.

Upon successful authentication, the user is redirected to a homepage on the IAManager. Through the homepage an authorized user can create events and begin interacting with wide-band displays. SAGE2 users are broadly categorized into three access levels in this new access control paradigm: administrator, event owners, and event participants. The administrator is responsible for granting event creation privileges to other users in the system apart from managing the different SAGE2 instances connected to the IAManager. In this model, events are the only way for users to access wide-band displays. We utilize the ABAC model to make an event a well-defined entity.

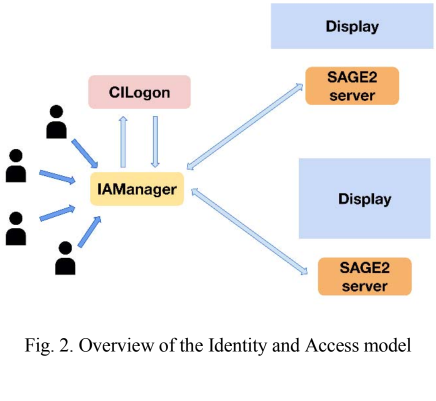

Fig. 2. Overview of the Identity and Access model

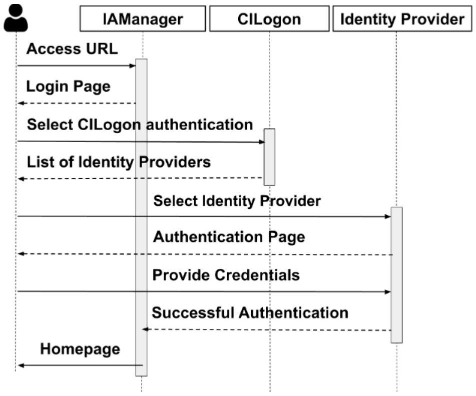

Fig. 3. Authentication using CILogon

An event has a schedule. That is an event has a start date, an end date, how often it recurs, meet time, and duration of every session. Events can be of three different types based on how open the event is. An event can be a private event having a closed group of members who have access to content created and shared as part of the event. The membership in this case is by invitation only. The event owner will add the members to the event. This type of event caters to groups that work with sensitive information that needs to remain hidden from anyone outside the group. For instance, a team working on a patent idea might need to safeguard its work from anyone who is not a member of the team. The second type of event is a little relaxed on who can be a member, in that anyone within the same organization can join the event without needing an explicit invite from the owner, although the owner will be able to manage the membership and remove any member the owner deems should not be a part of the event. Groups might have an open meeting policy allowing anyone within the organization to drop by at any time when the meeting is in session. Similarly in Universities, classes might be open to any registered university student. In such cases, privacy of data is not as much of a concern as accountability is. That is, the focus is on keeping track who is a part of the group and who is doing what within the group. The third type of event is open to the public. Anyone with a URL to the event can join the event. This type of event is for groups that do not have secure access requirements on the content that is generated or shared during the group meetings. This type of event is meant to accommodate cross organizational teams, student meeting, demos, and so on. The main utility of this type of event is the encapsulation it provides for the group and the content it generates. It must be noted that while the criterion for membership is increasingly more open in the second and third type of events, the owner of such an event may choose to close the membership of the event when they deem that their group has all the intended participants in it, akin to closing the door after the intended participants are seated inside a room. For instance, a faculty member might choose to close membership of the event for a particular course after the initial "add/drop" period for that course is over. A "default" event which is of the third type will be active when no other event is in session. This is to allow quick access to the wide-band displays without having the explicit need to create an event. This default event is useful mainly in enabling one-time meetings and similar situations where the utility of the wide-band display is more in terms of the screen real estate and ease of sharing content on the display rather than secure access to stored media and other CI. Any content generated during this default event is publicly accessible and stays on the SAGE2 server until explicitly removed. This default event is exactly similar in its behavior to the way a SAGE2 instance worked prior to the implementation of the IAM model.

Members of an event can have different roles assigned to them and this assignment can be modified anytime during the event. Every event is associated with a specific wide-band display (instance of SAGE2). A user can create events on a wide-band display after obtaining the requisite access permissions for creating an event on that display. There are two ways in which a user gets event creation permissions: (a) Through the IAMnager user interface homepage, a user can request the administrator for event creation permissions for any wide-band display. (b) The administrator can include a list of users in a configuration file to grant them event creation permissions. This allows someone at a managerial level to gain access to wide-band displays that they usually interact with, without having to go through the explicit request-approval process whereas a student, or employee for instance, who wants to create an event on a wide-band display will be able to do so after obtaining permission. Event owners can add users to the event, create custom roles for the event and assign or change roles of members. An event can be created on any server linked to the IAManager. This is usually a list of all the wide-band displays within the organization. Fig. 4 shows the architecture of the IAM system, detailing how it interconnects with CILogon for authentication and how it acts as a portal to accessing SAGE2. IAManager allows its users to create events on SAGE2 servers and add other users to the event. Whenever an event is created, IAManager stores the event details (such as when the event is scheduled, who the participants are, and so on) in a database. These details are fetched by SAGE2 servers at the start of an event. The event scheduler module of the SAGE2 server is responsible for scheduling of events including context switching and authorization of event participants.

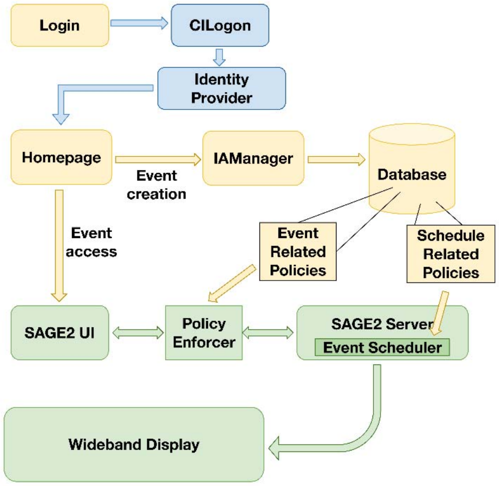

Fig. 4. Architecture of the IAM system

At a basic level, access control in SAGE2 translates to controlling the actions a user can perform while interacting with SAGE2. Since this new model restricts users to interacting with SAGE2 only through events, it makes way for discretionary access control, where the event owners can decide what interactions the rest of the members of the event are able to perform. Towards this end, an event owner can create roles for the event. A role is a collection of permissions that is identified by a name. When creating a role for an event the owner of the event will select a subset of SAGE2 interactions that a user with the role will be able to perform. Additionally, event level access control related permissions such as the ability to create new roles or changing role assignments of members can also be specified as part of the role description. This is particularly useful for creating event moderator roles, such a teaching assistant in a class event who will be responsible for managing the other participants in the event such as students. If U is the set of all interactions and M is the set of meta interactions such as defining a new role or changing the role assignment of a user, then a role R is simply a named subset of $(U \cup M)$ that an event owner can create as part of an event. An assignment of R to a user then ensures that the user is only able to perform interactions from the set R. An event owner can create any number of role definitions for an event and assign those roles to different users in the event. ABAC is responsible for enforcing the authorization of different interactions by users. Since this enforcement happens at runtime (shown in Fig. 4) by consulting the stored role definitions, both creation of new role definitions as well as (re)assignment of roles to different users can happen at any time during the event. The role information is passed to the SAGE2 server for which the event is created, and the access control related information of the role is stored with the IAManager.

A. *File Access Permissions*

File access refers to controlling who can view or modify a file. File access permissions defined on a file serves as rules in defining the access for a file. These permissions provide a secure way to manage the *assets* component of the SAGE2 access model. Access in SAGE2 is further refined by having access permissions on the content that is generated in an event. Every file that is created in an event will have two levels of access, that is, private to event, and publicly accessible. Content that is private to an event will be accessible only by members of the event. Access to event content translates to both putting it on the wide-band display and interacting with it when the event is in session as well as accessing the content offline when the event is not in session. The former kind of access is granted to all members of the event whereas by default only the participant who created the file has the latter kind of access. The owner of the file may choose to give offline access permissions to anyone in the event. Default access settings on content generated during an event is 'Private to Event'. This also serves to define the lifespan of the content. Anything marked private to the event gets purged when the event expires, and public files stay on the SAGE2 server until explicitly deleted. When the event expires, the event owner is prompted with an offline notification that allows the event owner to save a copy of the contents before they are purged.

Even though we have three types of events based on who can be a member, we only have two access levels defined for any asset generated in an event. Since an event of the second type will already have anyone from within the organization as a member, thereby granting them access to the content generated in the event, we decided to limit the access levels of assets to private and public since we feel that the permutations arising from these event types and access levels cover most of the group access scenarios. Furthermore, private assets of events that are open to organization or public is best explained through the example of a classroom. A course might be open for anyone within the university, however once the course registration ends, the course instructor might close the membership of the event. In this case, a student who is not registered for the course shouldn't be able to access the content, especially since accessing mostly translates to interacting with the content on the wide-band display.

B. *Context Switching*

Context of an event refers to the resources and applications that a user can access during an event. This includes open documents and application windows on the wideband display (including those shared by remote sites) as well as files previously generated as part of the event and stored on the SAGE2 server. Context switching is the process of replacing the context of one event with that of another. File permissions are used to filter the view of the media browser shown to users so that only files of that event are visible to users of an event along with any files from other events that are marked public. This ensures that users of one event can only access the media pertaining to that event. We thus encapsulate a team of users and the content they generate within an event. Further, using the ABAC model, event scheduling is carried out through a context switch mechanism built into the SAGE2 server. A SAGE2 server maintains a list of all the events that are created on it. As part of an event's details, the event schedule is known to the SAGE2 server. When it is time for an event session to start, the SAGE2 server makes a context switch (shown in Fig. 5) where it replaces the context of the previous event with the context of the current event. Such a context switch entails two things: (a) Updating the user media browser view with the list of files for the current event: Any users currently accessing the previous event will be notified of the end of the event session and unless the user is also a member of the current event, the user will be exited from the user interface of the wide-band display and will no longer be able to access the event until it is in session the next time. For the duration of the current event session, any user accessing the user interface of the wide-band display will see the updated view of the media browser showing the files of the current event only. (b) Loading the "state" of the current event onto the wide-band display: When an event session ends, the SAGE2 server saves the current state of the wide-band display as part of the context of the event that just ends and clears the wide-band display. The state of the wide-band display includes different application windows that are open on the wide-band display, their positions and sizes, and any other application specific states of individual apps. When a new event is about to start, the SAGE2 server will restore the previously saved state of the wide-band display for that event. This way, event members can continue their work from where they left it at the end of the previous session.

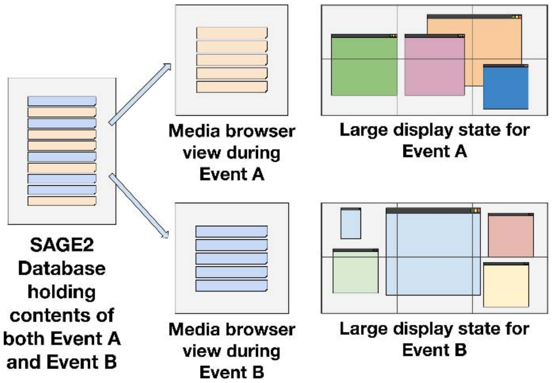
SAGE2 Database holding contents of both Event A and Event B
Media browser view during Event A
Large display state for Event A
Media browser view during Event B
Large display state for Event B

Fig. 5. Context Switch

V. Implementation

The IAManager has been implemented using Express, a Node.js web application framework. An npm module called "openid-client" which implements a server side relying party for Node.js runtime has been used to integrate with the CILogon's OIDC interface. CILogon requires that the relying parties be registered as clients with its server to ensure the authentication requests are coming from only the registered clients and to know what url to redirect the users to, once they are authenticated. The relying party client credentials are then used with every authentication request to verify the requests. The front end user interface for the IAManager uses React.js [15]. As shown in Fig. 6, after logging in the user sees a dashboard with all the events the user has created or has access to. The owner of an event is able to create new roles for each event or change role assignments of other users using the event menu. Fig. 7 shows a form that is presented to the users when a new event needs to be created. Here, the user can define the general behavior of the event including specifying what display wall the event will be run on. We make use of role-acl [16], an npm module to implement the ABAC model on the SAGE2 server. Whenever a new event is created a new set of policies, that define the event, is created and stored as a JSON object in the database. This JSON object contains information about (1) all the roles that are a part of the event and (2) a mapping of all the users of the event to the roles. At the beginning of every event session the corresponding JSON object is relayed to the SAGE2 server. This JSON object is used by the role-acl module to "grant" (shown as "Policy Enforcer" in Fig. 4) access permissions to the users of the event, based on their role definitions. Based on the role definitions the SAGE2 user interface of every user gets customized at run time to enable only those interactions to the user, that is defined by the user's role.

VI. Added Benefits

A. *Securing the Display Client*

SAGE2 transforms a wide-band display into a collaborative workspace through a display client that is open on the wide-band display and a web user interface client that every SAGE2 user opens on their personal devices. Limiting access of the user interface client to authorized users will ensure that only legitimate users are able to interact with the wide-band display. However, that still leaves the display client unprotected. That is, anybody would be able to open a display client on their device and see the contents on the display even though they may not be able to interact with it. The solution then is to place the display client under a similar authorized access as with the user interface client, but since the display client is intended to be run on the wide-band display which is typically a different device than the ones that are used to interact with SAGE2, the display client and the display system are not equipped with proper user interface to handle user authentication and authorization. The device authorization flow of the OpenID connect is best suited to handle this problem. Through the device authorization flow the display client authorization will happen by redirecting the authentication to a pre-registered personal device of the person managing the SAGE2 installation. That is, whenever a display client tries to connect to the SAGE2 server, a verification dialogue will prompt the SAGE2 admin on a chosen device (separate from the display device) such as a smartphone to enter the credentials. This way, only authorized display clients will be able to connect to the SAGE2 server.

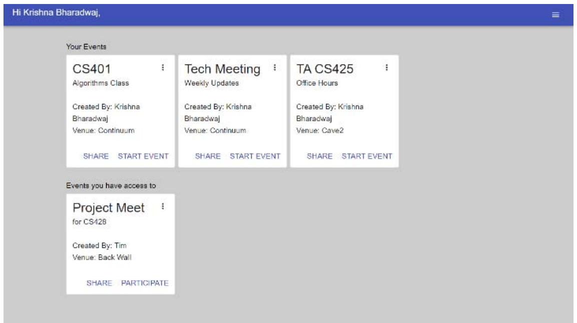
Hi Krishna Bharadwaj,

Your Events

CS401
Algorithms Class
Created By: Krishna Bharadwaj
Venue: Continuum

Tech Meeting
Weekly Updates
Created By: Krishna Bharadwaj
Venue: Continuum

TA CS425
Office Hours
Created By: Krishna Bharadwaj
Venue: Cave2

Events you have access to

Project Meet
for CS428
Created By: Tim
Venue: Back Wall

Fig. 6. User Dashboard showing Events

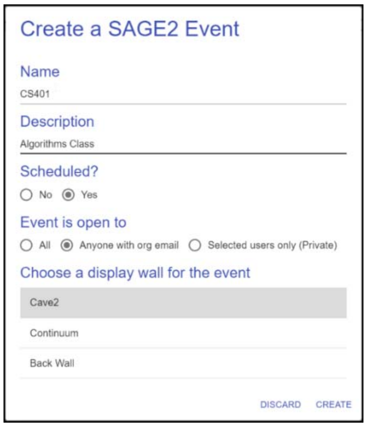
Create a SAGE2 Event

Name
CS401

Description
Algorithms Class

Scheduled?
○ No ● Yes

Event is open to
○ All ● Anyone with org email ○ Selected users only (Private)

Choose a display wall for the event
Cave2
Continuum
Back Wall

Fig. 7: Event Creation Form

B. *Remote Access*

Collaborators at one SAGE2 display site commonly work with remote colleagues who have their own SAGE2 display wall. Currently, during such remote collaboration, users can share documents with the remote collaborators and replicate their pointers on the remote SAGE2 display wall. Different sites may have specialized resources (for example, an image processing service) and enabling remote users to avail such specialized resources will enhance the remote collaboration. Remote access refers to giving users permission to access resources on a remote SAGE2 server. Taking advantage of the Single Sign On form of authentication, the tokens obtained from the OpenID connect authentication can be passed to remote SAGE2 servers and other protected resources for a fine-grained access control of resources on the remote machine. This can be particularly useful for addressing the issue of giving access to resources to a user of a remote SAGE2 server who might want access during a remote collaboration session between two SAGE2 wide-band displays. That is, even when a (remote) user is not an authenticated user of the home organization, as long as the user is authenticated by the remote organization and is part of an event which involves a remote collaboration between the two SAGE2 sites, the user will still be able to access the necessary resources as part of the event without having to repeat the authentication process at the home organization.

C. *Time Allocation on a Wide-band Display*

Events give rise to the possibility of time allocation. With the new IAM model in place, events can provide a secure way of interleaving the work of different groups to optimize the wide-band display utilization. The context switching makes way for a frustration free scheduling of works of different groups on the wide-band displays.

VII. Evaluation

With the identity and access management in place, SAGE2 allows for a more robust and secure collaboration using wide-band displays. As a preliminary evaluation we compare some of the key aspects of secure collaboration of IAM enabled SAGE2 with that of some of the popular collaborative tools available. Table 1 shows a comparison of the new SAGE2 model of identity and access management with a few popular collaborative platforms. Note that we have presented Google Meet and Google Drive as two separate entities in this comparison simply because they continue to be used for collaboration independently of one another. The table also maps different collaboration aspects to the components of the SAGE2 access model: users (U), physical display walls (W), meetings (M), assets (A), and cyber-services (C). As can be seen from Table 1, IAM enabled SAGE2 incorporates several access control related aspects that are available across the popular collaboration tools. Microsoft Teams, Google Meet, and Zoom support collaboration beyond video conferencing by integrating whiteboard functionality into the video conferencing. They have incorporated the notion of roles to decide who gets to edit the whiteboard at any given point, giving the ability to the meeting owner to toggle the whiteboard between presenter only and free-for-all modes. Essentially this means that the roles of participants can change as needed during a meeting. However, Teams, Meet, and Zoom all have a small, fixed set of roles just enough to enable presenter and free-for-all modes. SAGE2's support for role-based access similarly allows role assignments to be dynamically changed and in addition allows creation of custom roles that are tailored to the needs of individual groups.

TABLE 1: COMPARISON OF SAGE2 IAM WITH POPULAR COLLABORATION TOOLS

Secure Collaboration Aspects	**SAGE2 Access Model Components**	**SAGE2**	**Google Drive / Docs**	**Zoom**	**Microsoft Teams**	**Google Meet**	**Webex**	**Whereby**
User Login	U	✓	✓	✓	✓	✓	✓	✓
File Creation and Storage	A, C	✓	✓	x	✓	x	✓	x
Document Sharing	A	✓	✓	✓	✓	✓	✓	x
Private Folders and Workspace	U, A	x	✓	x	✓	x	✓	x
Role Based Access Support	U, A, C	✓	✓	✓	✓	✓	✓	x
On-Screen Access Control	U, A, C	✓	✓	✓	✓	✓	✓	x
Custom roles within events / meetings	U, A, C	✓	x	x	x	x	x	x
Secure Meetings	U, A, M, C	✓	x	✓	✓	✓	✓	✓

VIII. Conclusion and Future Work

This paper describes (1) User requirements with respect to security concerns in wide-band displays. (2) Some of the main access control models in use. (3) A novel access control model for providing secure access to wide-band displays for group work (5) Identity and Access Management System based on the proposed model and its implementation details. (6) A few added benefits that come with the implementation of the proposed model. (7) An evaluation of the system.

By defining three different types of events with varying levels of security, our proposed model allows for groups with varying security needs to share the same platform for carrying out their group work. Furthermore, by allowing custom roles to be defined for each event, our model supports groups with different group dynamics. Thus, event owners can create roles that are more meaningful and relevant to each group. Also, our proposed model allows group administrators and moderators to easily manage group membership and dynamically control access within the groups by changing role assignments without the need of security administrator's intervention. By enabling custom roles to be created for each event and enabling role assignments to change dynamically, the SAGE2 IAM model takes into account the possibility of group dynamics evolving with the passage of time and the possibility that users may don multiple hats during the course of an event.

In the future, the SAGE team plans to conduct user evaluation of this new security system to see how effectively it meets the user requirements that were detailed in this paper and also to explore how this new model impacts the various SAGE2 usage patterns [7].

References

[1] Mary Czerwinski, George Robertson, Brian Meyers, Greg Smith, Daniel Robbins, and Desney Tan. 2006. Large display research overview. In CHI '06 Extended Abstracts on Human Factors in Computing Systems (CHI EA '06). Association for Computing Machinery, New York, NY, USA, 69–74. doi:https://doi.org/10.1145/1125451.1125471

[2] X. Bi and R. Balakrishnan. "Comparing usage of a large high-resolution display to single or dual desktop displays for daily work". In Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '09, page 1005–1014, New York, NY, USA, 2009. Association for Computing Machinery. ISBN 9781605582467. doi:https://doi.org/10.1145/1518701.1518855

[3] B. Jeong, L. Renambot, R. Jagodic, R. Singh, J. Aguilera, A. Johnson, J. Leigh. "High-Performance Dynamic Graphics Streaming for Scalable Adaptive Graphics Environment", SC '06: Proceedings of the 2006 ACM/IEEE Conference on Supercomputing, Tampa, FL, USA, 2006, pp. 24-24, doi: 10.1109/SC.2006.35.

[4] T. Marrinan, J. Aurisano, A. Nishimoto, K. Bharadwaj, V. Mateevitsi, L. Renambot, L. Long, A. Johnson, J. Leigh. "SAGE2: A New Approach for Data Intensive Collaboration Using Scalable Resolution Shared Displays," In Proceedings of the 10th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing, Miami, FL, October 22, 2014.

[5] L. Renambot, T. Marrinan, J. Aurisano, A. Nishimoto, V. Mateevitsi, K. Bharadwaj, L. Long, A. Johnson, M. Brown, J. Leigh. "SAGE2: a Collaboration Portal for Scalable Resolution Displays," Future Generation Computer Systems, Volume 54, January 2016, Pages 296–305, doi:10.1016/j.future.2015.05.014

[6] J. Leigh, A. Johnson, L. Renambot, T. Peterka, B. Jeong, D. J. Sandin, J. Talandis, R. Jagodic, S. Nam, H. Hur, Y. Sun. "Scalable Resolution Display Walls," Proceedings of the IEEE, Issue 99, May 2012, pp. 1-15, doi:10.1109/JPROC.2012.219160.

[7] J. Leigh, M. Belcaid, D. Kobayashi, N. Kirshenbaum, T. Wooton, A. Gonzalez, L. Renambot, A. Johnson, M. Brown, A. Burks, K. Bharadwaj, A. Nishimoto, L. Long, J. Haga, R. Pelayo, J. Burns, F. Cristobal, J. McLean. "Usage Patterns of Wideband Display Environments In e-Science Research, Development and Training", eScience 2019, San Diego, California, USA, September 24-27, 2019

[8] D. R. Kuhn, E. J. Coyne and T. R. Weil. "Adding Attributes to Role-Based Access Control", in Computer, vol. 43, no. 6, pp. 79-81, June 2010, doi: 10.1109/MC.2010.155.

[9] D. F. Ferraiolo and D. R. Kuhn. "Role-Based Access Controls" in Proceedings of the 15th National Computer Security Conference, Baltimore, Oct 13-16, 1992. pp. 554 - 563

[10] J. Basney, H. Flanagan, T. Fleury, J. Gaynor, S. Koranda, and B. Oshrin. "CILogon: Enabling Federated Identity and Access Management for Scientific Collaborations". In Proceedings of the International Symposium on Grids and Clouds (ISGC), PoS(ISGC2019)031, 2019. doi:https://doi.org/10.22323/1.351.0031

[11] https://openid.net/connect/

[12] Alfieri R. et al., VOMS, an Authorization System for Virtual Organizations. In: Fernández Rivera F., Bubak M., Gómez Tato A., Doallo R. (eds) Grid Computing. AxGrids 2003. Lecture Notes in Computer Science, vol 2970. Springer, Berlin, Heidelberg. doi:https://doi.org/10.1007/978-3-540-24689-3_5

[13] Ceccanti, A., Hardt, M., Wegh, B., Millar, A., Caberletti, M., Vianello,E.,Licehammer,S., The INDIGO-datacloud authentication and authorization infrastructure. Journal of Physics: Conference Series 898, 2016. doi:https://doi.org/10.1088/1742-6596/898/10/102016

[14] M. Uddin, S. Islam and A. Al-Nemrat. "A Dynamic Access Control Model Using Authorising Workflow and Task-Role-Based Access Control," in IEEE Access, vol. 7, pp. 166676-166689, 2019, doi: 10.1109/ACCESS.2019.2947377.

[15] https://reactjs.org/

[16] https://github.com/tensult/role-acl#readme

[17] US Department of Defense, Trusted Computer System Evaluation Criteria. In: US Department of Defense (eds) The 'Orange Book' Series. Palgrave Macmillan, London, 1985.

[18] H. Ould-Slimane, M. Bande and H. Boucheneb. "WiseShare: A collaborative environment for knowledge sharing governed by ABAC policies," 8th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom), Pittsburgh, PA, USA, 2012, pp. 21-29, doi: 10.4108/icst.collaboratecom.2012.250402.

[19] P. Nasirifard, V. Peristeras, and S. Decker. "An annotation-based access control model and tools for collaborative information spaces". Emerging Technologies and Information Systems for the Knowledge Society, pages 51-60, 2008.