Editor's Note

In the summer of 2024, Clifford Lynch announced his retirement as executive director of the Coalition for Networked Information (CNI) after 28 years at its helm. CNI quietly launched a project to create this Festschrift to document and honor his legacy. Authors began contributing articles in early 2025, with a planned publication date of July 2025. Since the final membership meeting of Cliff's tenure was April 7–8 in Milwaukee, the plan was to surprise him, surrounded by colleagues and friends, with a presentation of the table of contents of this special issue. However, just two weeks prior to the meeting, Cliff's health worsened; he was told about the Festschrift and received project details and articles. Though unable to attend in person, he participated in the CNI membership meeting via Zoom and also virtually joined his retirement reception, which included readings of excerpts from each article in this volume. Sadly, on April 10, 2025, Clifford Lynch passed away. Festschrift contributors wrote their articles prior to his passing, and we have chosen not to alter their original language.

An Interview That Ran Long

Ken Klingenstein

abstract: From 1986 to the present, I have had the fortune to have wide-ranging conversations with Cliff Lynch on networks and information. Beginning with discussions of networking requirements for libraries, our exchanges quickly encompassed such topics as access control, security, and privacy. Increasingly we focused on aspects of middleware, a layer of services above the network but below applications, including developing a framework for user identifiers and mechanisms for authentication. Cliff's contribution to the conceptualization of "edge" cases, the unique cases that higher education presents, was particularly valuable in these developments. Over the years, thousands of research universities, companies, and governments around the globe have benefited from the infrastructures built, in part influenced by Cliff's unmatched perspectives. In addition, he encouraged discussion of this important set of topics at Coalition for Networked Information meetings, broadening participation and dissemination.

Introduction

In 1986, I was invited to visit the University of California, Berkeley, to review their campus network. One of the numerous people there who needed to be interviewed was a young librarian who was said to have a good sense of the future. My prepared questions were oriented toward the future networking requirements for libraries. The interviewee, Cliff Lynch, took the conversations to several other places, however, including topics of privacy, access control, and security. Our conversation wandered this landscape far longer than the time set aside, leaving the rest of my site visit schedule in shambles. At the end, Cliff and I agreed that we would find some way to continue the "interview." For almost 40 years, we have.

The Emergence of Middleware

As the Internet began to develop, in the mid 1980s, it became clear that a layer of services above the network but below the applications would be needed to scale connecting users to applications. The initial motivating applications were all research uses that highlighted advanced network capabilities but lacked any mechanisms to support the larger num-

portal: Libraries and the Academy, Vol. 25, No. 3 Supplement (2025), pp. 93–98.

bers of users that were about to come. (The advanced applications were called "network stunts" by their developers, recognizing that actual wide deployments would require something in the middle to scale a user base.) Users would need digital identities, which would have to provide a way to authenticate that a network request really came from them. Those now-authenticated users would need attributes about themselves—that is, information that characterized the person—to be provided to the applications. All those mechanisms had security and privacy dimensions. The term *middleware* was coined to describe this layer of connectivity above the network but below the application layer.

In middleware, particularly for identity and trust, use cases are the critical seed corn. They are well chewed, structured, prioritized, and then used to drive an architecture that can meet them. Cliff has been, and continues to be, the encyclopedia of use cases. When Cliff says that something is important, I pay attention.

When Cliff says that something is important, I pay attention.

My conversations with Cliff in those early years were largely about the Internet and how it might evolve from an academic invention to a wider world. Even then, long before middleware would be crafted, we would consider the key issues of trust and identity in a networked world, discussing approaches that might preserve privacy. Though the community of users was still small, we knew that the technologies were mighty and would need to scale globally. Discussions such as indirection for identifiers to preserve privacy and what mechanisms might resolve that indirection had begun to happen, and Cliff provided critical perspectives as we formulated these issues. And, as always through many years of interviews, Cliff's knowledge of what was happening in the digital content and services world was all-encompassing. I never left a dialog with him without a list of things to look up.

. . . Cliff's knowledge of what was happening in the digital content and services world was all-encompassing.

Authentication and Identity

One central issue for us budding middleware activists in those days was to refactor emerging networked information applications to externalize some of their required functionalities, particularly authentication and identity. It was a tough sell, given that the application developers were most interested in the cool features of their apps. They were reluctant to spend energy or funds on mundane if essential underlying services like authentication. Moreover, even if they wanted to externalize some services, there was at best a nascent institutional infrastructure for them to leverage. It was a classic chicken and egg situation that Cliff helped resolve. He promoted both sides of the issue effectively, encouraging the app developers to think about externalizing identity and urging the institutions to build the infrastructure that those applications would leverage. The results have been remarkable. Well over a thousand universities and research centers nationally, and tens of thousands internationally, as well as corporate and government entities now use that identity trust infrastructure. Cliff was catalytic in those efforts.

The Growth of Internet2

In the early 1990s, the middleware layer began to receive concerted attention. In 1996, the U.S. research and education community started a project known as Internet2, originally focusing on connectivity, speed, and reliability.[1] Today, it is a nonprofit research and development consortium led by hundreds of U.S. universities in partnership with industry and government, and its focus now includes trust and identity.[2] Internet2 and EDUCAUSE became partners in this work, with Internet2 doing most of the development work. It utilized a series of National Science Foundation (NSF) grants to develop middleware software such as Shibboleth, which allows people to sign in using just one identity to systems run by different universities or institutions,[3] and Grouper, which enables users to create and manage institutional and personal groups.[4] Internet2 also developed schema standards such as eduPerson, which stores attributes about a user, such as employee or student status or graduating class.[5] EDUCAUSE created a series of activities to educate campus IT leaders on approaches to deploying the software and developing the business processes to provide users with accounts and attributes. The Association of Research Libraries (ARL) and the Coalition for Networked Information (CNI) engaged the library community; National Information Standards Organization (NISO) began to adopt some of the products of the work as standards.

The activities also expanded internationally. Jisc in the United Kingdom (UK) was quick to engage. GÉANT (Gigabit European Academic Network), the European Union (EU) network roughly equivalent to Internet2, added a portfolio of middleware efforts. Due to the multinational character of the EU, GÉANT became an early proponent of the multilateral approaches being pioneered in the United States. These efforts culminated in a meeting hosted by the United Kingdom in Lower Slaughter in the Cotswolds in 2004, the first real gathering of national trust and identity federations. (The meeting had the tagline "Leading Trust to Slaughter.")

In the 1990s, as Cliff and I continued our periodic exchanges, we started to talk more about the basics of identifiers and authentication. There are important characteristics of identifiers, whether they apply to people, objects, or content, that needed to be enumerated to make sure that those crafting standard identifiers understood the consequence of their decisions. Persistence, opacity, embedded metadata, resolving identifiers, and many other considerations began to manifest themselves in these conversations. In the middleware space at the time, people were the most important objects to identify, but Cliff was already thinking about how these identifier characteristics applied to content and then to derived content. Intuiting taxonomies has always been one of Cliff's strengths.

Intuiting taxonomies has always been one of Cliff's strengths.

Around that time, Paul Evan Peters, the founder of CNI, died suddenly. Paul was a distinctive genius, yet gracefully, CNI preserved its unique role in the scholarly ecosystem by selecting Cliff to direct it. The charisma of Paul's booming voice was followed by the equally compelling charisma of a Cliff Lynch mumbled aside.

Access Control

At the turn of the century, while the rest of the IT world was chasing the year 2000 problem (Y2K), Cliff and I began our first discussions of attributes for access control. While large-scale attribute controls (such as permissions for library walk-in patrons) were our initial use cases, the "Identerati" were most interested in the "edge" cases. Simple designs might handle the bulk of cases and might be addressed by the marketplace, but Higher Ed-ge, as we called our use cases, had unique requirements that needed more fine-grain controls. One of the best places to look for these was in special collections in libraries, which often have distinctive access requirements. Controls might be by domain, or accreditation, or copyright restrictions. Cliff's knowledge base was invaluable in supplying those edge cases to us, and the architectures that those use cases inspired proved extremely durable over time. My interviews with Cliff during that period continued to run long.

In the 2000s, the community discussions moved into the deep waters of authentication. It was not a particularly interesting set of issues, but our primary customer for trust and identity was the federal government, and for them, it was mission critical. Identity proofing and the technologies of authentication took front stage and had little consideration of privacy. It proved to be a set of issues where even the ratholes had ratholes. In the midst of those long labors, some of us continued to assert that authentication had to be integrated into traditional library values, including the freedom to roam the shelves in an academic library anonymously. Identifying those values and selecting technology choices that preserved those values became critical decisions, and Cliff's continued input in those discussions was invaluable.

Attributes and Scholarly Software

The higher education community also took a first deep dive into attributes in the 2000s, leading to the Tao of Attributes workshop at the National Institutes of Health (NIH) in 2009 (https://videocast.nih.gov/watch=8057). It was truly a "throw it long" discussion, but a welcome relief from the grind of details that are intrinsic to authentication. The spectrum of open questions ranged from how to craft attributes to how to use them and how to name them. Another matter discussed was how to control attribute vocabulary—that is, the values an attribute could take—so that enough but not too much meaning was conveyed. If authentication was a slog in details, attributes were green fields ripe for invention. Over the next few years, Cliff was a key guide, offering invaluable suggestions as the trust and identity crowd began to frame what work to do and when. It has been a continuing personal pleasure to watch Cliff graze in our green fields.

The Mellon Foundation, under the vision of Don Waters, was making important decisions about supporting key open-source projects for scholarly software. Repositories, on-line learning systems, digital publications, and reference taxonomies were emerging concepts that needed implementations. The community did not know enough to understand what features would be important, so it was appropriate to let several flowers bloom and then help them cross-pollinate to produce superior products. Cliff was one of those vital cross-pollinators, a busy bee assisting such software projects as

DSpace, Fedora, VIVO, and Sakai to better themselves. With universities focusing on identity management infrastructure and with federal agencies, notably NSF, investing in networked research software, there was a gap in resources to create scholarly software. The Mellon Foundation, with wisdom frequently offered by Cliff, helped fill that gap, and many of those projects have reached maturity and sustainability.

Persistence and Preservation

In the 2010s, we began to wrestle with persistence and preservation. The sheer volume of materials available, and the lack of normative taxonomies, made discovery impossible. Even the persistent identifiers turned out not truly persistent, and sustainable mechanisms had to be set up to insure their permanence. Cliff's advocacy of the needs and his endorsement of the solutions were important. We needed ORCIDs (open researcher and contributor IDs), OIDs (object identifiers), URNs (uniform resource names), URIs (uniform resource identifiers), and more to add durability to our exploding networked information world. For all of us working in this space, it was daunting to stay informed about all the activities taking place. Cliff threaded a fine needle in his service to the community. Recognizing the resource limits of CNI and the networked information community in general, Cliff stepped into the unique role of trusted information broker, both as an individual and in making it the central programmatic goal of the organization. With his large network of colleagues and confidants, he could sift through what to communicate to whom.

Other Issues

Another emergent trend in the teens of the twenty-first century was multiple personas and devices per user. Understanding the complexity of what parts of which personas to place on which devices is a multidimensional space, one where Cliff's eclectic insights were important. Moreover, the "ownership" of personal information passes from the user to the service. Cliff played an important role in keeping that concern visible.

As social media emerged, the scholarly issues around studying its impacts became another focal point of conversation. Referencing, archiving, patterns of use, and other dimensions became critical areas of study, especially as social media grew critical within our society. With Cliff at CNI, there was a place for these important discussions to take place.

Finally, our interviews recently acknowledge the ground we have covered in 40 years, so we have begun to go "meta" about what we have seen. We had witnessed enough middleware activities that we could begin to talk about strategies for building technologies and the communities that use them, analyzing what has worked and what has not. Few people have the sweep of the societal and scholarly landscape that Cliff has. His perspectives are unmatched.

Few people have the sweep of the societal and scholarly landscape that Cliff has. His perspectives are unmatched.

Notes

1. Frank DeMaria, "The 20-Year History behind Internet2," *Campus Technology* (October 22, 2016), https://campustechnology.com/articles/2016/09/22/the-20-year-history-behind-internet2.aspx.
2. Internet2, "About Us," 2025, https://internet2.edu.
3. Shibboleth Consortium, "The Shibboleth Project," 2025, https://www.shibboleth.net/about-us/the-shibboleth-project; Barbara Losoff, "Shibboleth: A Project of the Internet2 Middleware Initiative," *Collaborative Librarianship* 1, 1 (2009), https://digitalcommons.du.edu/collaborativelibrarianship/vol1/iss1/3.
4. InCommon, "Grouper," 2025, https://incommon.org/software/grouper/; Illinois State University, "What Is Grouper?" 2024, https://help.illinoisstate.edu/identity/guides/what-is-grouper.
5. Texas A&M University Technology Services, "Eduperson Schema Attributes," https://docs.security.tamu.edu/docs/identity-security/attribute-services/attribute/edupersonschema/.