CENIC Administrative Policy and Practice (CAPP 16)
CENIC Network Data Privacy Policy

As Revised on July 21, 2016

I. Introduction

The Corporation for Education Network Initiatives in California (CENIC) is a nonprofit corporation formed in 1997 to provide high-performance, high-bandwidth networking services to California universities and research institutions. CENIC operates the California Research and Education Network (CalREN), a high-capacity network designed to meet the unique requirements of over 20 million users, including the vast majority of K-20 students together with educators, researchers and other vital public-serving institutions. By interconnecting all of these organizations with each other, including sites with unique supercomputer centers and scientific user facilities, as well as a wide range of educational technology, millions of students, researchers, and members of the public are able to transfer data, access remote resources, and collaborate productively.

The CalREN backbone consists of roughly 3,800 miles of CENIC-owned and managed fiber plus last-mile fiber, hundreds of last mile circuits, customer-premises equipment (CPE), and thousands of optical components.

Three networks comprise CalREN:

CalREN-DC is CENIC's "Digital California" network, providing high-quality network services for students, faculty, researchers, and staff at education institutions. This network provides connectivity to the commercial Internet.

CalREN-HPR is CENIC's "High-Performance Research" network, and provides leading-edge services for large-application users at CENIC Associates sites. This network provides connectivity to Internet2 and ESnet.

CalREN-XD is CENIC's "eXperimental/Developmental" network, and consists of a set of build-to-order resources to support bleeding-edge services for network researchers at sites such as the San Diego Supercomputer Center, the University of California Institutes for Science and Innovation, the Center for Advanced Computing Research at Caltech and its Jet Propulsion Laboratory, the University of Southern California and its Information Sciences Institute, Stanford University and the Stanford Linear Accelerator Center, national laboratories and other major network research entities that collaborate with these researchers in California.

II. Scope of this Policy

What This Document Is — the purpose of this document is to convey to individuals and organizations who use CENIC's networks, as well as other members of the public, in high-level terms, what network-related data CENIC collects, and why, and what CENIC does or does not do with that data. This document does not represent a change from prior practices, but merely seeks to publicly convey what is already done. It is intended to be a "living document," but it is also designed to be high-level enough that it does ideally not require annual revisions as technological changes occur.

What This Document Is Not — this document is not a policy and procedure manual. For the reasons indicated above, it does not contain or attempt to contain details about how every possible situation can and will be handled, because such details can change frequently as some new security or

privacy technique or tool is made available or adopted by CENIC. Policy and procedure manuals exist within CENIC that do cover these details. This document is also not a legal contract with its constituents. Those Memorandums of Understanding co-exist separately from this document.

Specifically, this policy identifies:

- the information CENIC collects about data transferred by its infrastructure;
- the justification for this collection;
- the ways in which this information may be used and disclosed;
- the way that data is retained; and
- the security measures adopted to prevent unauthorized access to this information.

Note: Like most research and education networks, CENIC already publishes information about the amount of data it transfers (bytes, packets, unicast, multicast, errors, etc.), and the speed at which data is transferred (bits per second, bps trend over time) openly, and in real time (http://cenic.org/network/real-time-mngmt-tools), and will continue publishing this information. This data is highly aggregated, and does not contain information about traffic flows specific to individual users. Thus, this policy does not apply to this "performance" and "utilization" data.

III. Key Principles

CENIC's data privacy policy strives to balance the privacy interests of the users and institutions whose data transits CENIC's networks, the scientific needs of network researchers, the educational needs of schools, and the operational needs of CENIC. We are committed to protecting privacy and informing interested parties about our policies and practices.

IV. Information That May Be Collected

Flow Data: CENIC captures, collects, and evaluates network monitoring data ("flow data") for operational, troubleshooting, capacity planning, and research purposes. This data helps to aid in operational support, capacity planning, forecasting, fault diagnosis and resolution, cybersecurity, and also to support research projects and usage reporting.

Network monitoring data consists of electronic records that concisely characterize network transmissions. The records include data commonly referred to as "packet headers." These headers include the routine information needed by network infrastructure to forward packets and include such fields as: source and destination IP addresses and port numbers, protocol type, bytes transferred, timestamps, and network interfaces transited.

Packet Contents: CENIC has the technical ability to collect data packet contents, but only does so rarely – either in the process of testing or troubleshooting network connectivity or performance, or in coordination with a CENIC site. Internal procedures exist to ensure this is done securely. CENIC takes the following actions when collecting packet contents:

1. Permission is obtained from the cognizant executive (e.g., CEO, COO, CTO)
2. Affected organizations are notified
3. Internal logs are maintained documenting what data is collected, the time period covered, who collected the data, why the data was collected

V. How Data Is Collected, Retained, and Protected

A. Data Collection

All network monitoring data collected is managed under the control of authorized CENIC employees and contractors only. All collected data is maintained in electronic form only and is never converted to hard copy.

B. Data Protection

CENIC takes appropriate steps to protect collected data from unauthorized access or disclosure. However, CENIC does not encrypt any data collected. Additionally, CENIC employs industry standard security measures, including physical, electronic, and procedural safeguards such as Non-Disclosure Agreements (NDAs), to protect against the disclosure, loss, misuse, and alteration of the information under our control.

C. Data Retention

Data collected for network monitor and control purposes has value to CENIC for operations, engineering, and administrative functions. Different forms of data have varying useful lifetimes, depending on the data type and internal use. The general rule is collected data is retained only as long as it has a business purpose. The specific retention period for any data type is given in a CENIC internal policy and is subject to revision from time to time. Regardless of the policy, data may be preserved beyond the nominal retention period upon legal demand or at the direction of the CENIC CEO or CENIC counsel to preserve.

D. Data Disposal

When data previously collected are no longer relevant, such records are cleared using industry best practice[1] data sanitization techniques.[2] Data are overwritten in place with a series of fixed and random bit patterns using software approved by CENIC for the computing platform.

VI. Disclosure of Data

CENIC is the steward of all network monitoring data it collects within its infrastructure. In general, CENIC does not disclose, give away, or sell its network monitoring data to any other organization, nor does it delegate its stewardship responsibility. Notwithstanding this non-disclosure principle, CENIC may share network monitoring data under the following circumstances.

A. Member Institution Electronic Access

CENIC customers may from time to time make a request for information about ***their own*** usage of CalREN. In such cases, CENIC will make reasonable attempts to provide views of network usage data that only include information that the particular site could reasonably have gathered on its own, by analyzing its network connections to CENIC. Further, CENIC will coordinate with a designated representative of the site before sharing such data.

In the event of a network or cybersecurity emergency affecting CENIC sites, customers, or the wider Internet, CENIC may release relevant customer data to the customer to enable the process of troubleshooting, analysis, or service restoration.

B. Network Providers

CENIC is part of a large community of Research and Education Networks (RENs) at various scales (international, national, regional, state). These include, for example, Internet2 and ESnet in the United States.

Other Research and Education Networks may, from time to time, make a request for information about the traffic they exchange with CENIC (that is, traffic that passes between the two organization's networks). If data sharing is authorized, CENIC will make reasonable attempts to provide views of data that only include information that the organization could reasonably have gathered on its own, by analyzing its own network connections to CENIC. We require all organizations provided with such data to enter into an agreement prohibiting them from sharing it further.

C. Researchers

Network researchers may from time to time make a written request for network monitoring data. Where permitted to do so by agreements with CENIC constituents, CENIC will use current accepted practices in the R&E community,[3] such as anonymizing the IP addresses in a prefix-preserving manner[4] and/or removing low-order bits, to de-identify the data released to researchers to ensure user privacy wherever possible.

If data sharing is authorized, researchers must subsequently agree to specific terms of use for the project in question, and enter into a written agreement. Researchers are not permitted to share this data with any party for any reason, unless authorized in writing by CENIC. Release of data that has not been de-identified is rare, and requires network researchers to submit a written security plan for review and assessment.

Institutions may request that data about their utilization of CENIC's network resources not be shared at all. Such requests must be made in writing by a properly authorized representative of the institution, and must be renewed on an annual basis.

D. Third Parties

For the purposes of securing the network and to analyze and resolve operational issues, CENIC may, by contract, involve third parties. This analysis may include, from time to time, raw packet captures (i.e., both packet headers and contents) in addition to network monitoring (flow) data. Both CENIC staff and such third parties are obligated to protect the data and use it only for the purposes identified. CENIC will assure that such information is managed and shared with third parties only within a defined contractual relationship, and bound by a written Non-Disclosure Agreement. CENIC shall provide written notification to any affected institution when such data is collected, the purpose of the collection, and the inclusive dates and times of such collection.

E. Legal Requests

If required by law and upon advice of legal counsel, and approval of the CENIC CEO, CENIC will comply with lawful requests to disclose network monitoring data. CENIC shall provide written notification to any affected institution when such data is collected, the purpose of the collection, and the inclusive dates and times of such collection, to the extent permitted by law.

F. Other Requests

Other requests for network monitoring data will be processed with the assistance of CENIC's legal counsel and the approval of the CENIC CEO.

VII. Notice for Updates and Changes to Policy

CENIC reserves the right to update this privacy policy at any time to reflect changes in the manner in which it deals with traffic, whether to comply with applicable regulations and self-regulatory standards, or otherwise. The Privacy Policy posted here will always be current. We encourage you to review this statement regularly.

VIII. Who to Contact if You Have Questions

If you have any questions about this privacy policy, please contact: privacy@cenic.org

---

[1] Best Practices for Data Destruction, Privacy Technical Assistance Center (http://ptac.ed.gov/document/best-practices-data-destruction)
[2] NIST Special Publication 800-88, Revision 1, Guidelines for Media Sanitization (http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-88r1.pdf)
[3] An example set of criteria includes the process used by CAIDA (http://www.caida.org/data/passive/passive_dataset_request.xml)
[4] An example of a tool that does this is Crypto-Pan (http://www.cc.gatech.edu/computing/Telecomm/projects/cryptopan/)