

USC-ISI Hosting Infrastructure

USC-ISI 2012

1. Description

- To operate as a PREDICT Data Host (DH), the LACREND project maintains computers at USC/ISI, USC/HPCC and CSU. To operate as a PREDICT Data Provider (DP), the LACREND project operates computers at USC/ISI, USC/HPCC, CSU, FRGP, and Keio University.

2. System Inventory

Hardware

- At USC/ISI as DH: primary dataset distribution is done on a webserver (a Dell 4-core Intel Xeon box running Fedora 17 Linux). Front-end web service on ISI's commodity web server (a Solaris Sun-280R running Solaris 2.9).
- At USC/HPCC as DH: we operate on USC's High Performance Computing Cluster (details at <http://www.usc.edu/hpcc>). Among supercomputers in an academic setting, HPCC's supercomputer cluster is the 10th fastest in the United States. Among all supercomputers in the world, it is ranked 166th. It claimed this distinction by achieving a benchmark in Fall 2012 of 149.9 teraflops, or 149.9 trillion floating-point calculations per second, on its 2,225-node, 10-gigabit backbone cluster. HPCC comprises a diverse mix of computing and data resources. Two Linux clusters constitute the principal computing resource, with 785 dual-core Intel Xeons and 1990 quad-core Opterons. In addition, HPCC has a central facility that provides more than 1.2 petabytes of combined disk storage and potential access to nearly a petabyte of tape storage. Of this facility, LACREND maintains 26TB of dedicated disk storage. These facilities are used for processing anonymized packet header data and are planned to be used for researcher-to-data remote access.
- At CSU as DH: we operate three capture points, a data archive and compute nodes to support researcher- to-data operations. All three capture points run at 10Gb/s and they continuously capture traffic data, anonymized it and archive it in flow form. We have been archiving data since 2008. The archive contains about 96TB of storage, with another approximately 40TB distributed over various machines. The compute node is a high-end machine with 32GB or RAM and 4TB or storage. We are in the process of adding 10 more computing machines.
- At USC/ISI and USC/ HPCC as DP for anonymized packet headers: LACREND does not directly operate any packet collection facilities, but USC/Information Technology Services operates LACREND hardware that provides anonymized packet headers. These facilities include three Dell packet capture boxes with Endace capture cards operating at 1Gb/s and 10Gb/s.
- At USC/ISI, CSU, and Keio University as DP for internet address space data collection: LACREND operates Internet Address probing boxes. These boxes vary, but are at least Dell, 4-core Intel Xeon or AMD Opterons, typically with 2 or more network interfaces. We currently do address probing from three sites: USC/ISI, CSU, and Keio University's Shonan Fujisawa Campus. We are planning to deploy new address probing boxes at additional locations over the course of LACREND's period of performance.

- At CSU and FRGP as DP for packet flows: LACREND operates three 10Gb/s capture points. These are equipped with high-end Endace cards that allow line speed capture and ns time precision. Custom LACREND software merges the two flow directions and anonymizes the traffic. Our custom software also manages a RAM disk for cases where line speed operations need to be performed.

Software

- All of these computers run CentOS, Fedora, or Ubuntu versions of Linux.
- Data collection (the DP role) at USC/ISI, USC/HPCC, CSU, and FRGP is done with the LANDER packet processing software, described in the paper: "Experiences with a Continuous Network Tracing Infrastructure" by Alefiya Hussain, Genevieve Bartlett, Yuri Pradkin, John Heidemann, Christos Papadopoulos, and Joseph Bannister, In Proceedings of the ACM SIGCOMM MineNet Workshop, pp. 185-190. Philadelphia, PA, USA, ACM. August, 2005. <<http://dx.doi.org/10.1145/1080173.1080181>>.
- Data distribution (the DH role), is done at USC/ISI using custom web-based software to provide secure access for specific datasets to researchers, handling of dataset requests, and auditing of the process.
- For some sensitive datasets, data "distribution" (the DH role), is done using researcher-to-data, where researchers are given accounts and allowed to log-in to LACREND-operated machines. This model is currently operational at CSU, and it is planned for use at USC/HPCC.

Backups, mirroring, and redundancy

- We back up all archival data at least once.
- Backups at USC/ISI are done to an external RAID operated by ISI operations staff. Both master copies and backups are all checksummed and checksums for one dataset is checked each day.
- Backups at USC/HPCC are done by their operations staff using off-line tape storage.
- Backups at CSU are done via external USB drives as needed. The archive runs RAID5 for robustness, and the external drives are used to backup all the data. The drives are then stored separately from the archive.

USC-ISI 2013

1. Description

- To operate as a PREDICT Data Host (DH), the LACREND project maintains computers at USC/ISI, USC/HPCC and CSU. To operate as a PREDICT Data Provider (DP), the LACREND project operates computers at USC/ISI, USC/HPCC, CSU, FRGP, and Keio University.

2. System Inventory

Hardware

- At USC/ISI as DH: primary dataset distribution is done on a webserver (a Dell 4-core Intel Xeon box running Fedora 19 Linux). Front-end web service on ISI's commodity web server (a Solaris Sun-280R running Solaris 2.9). Storage is handled by two storage servers. Three are Dell 2950s with fully populated internal disks running as JBOD, and with external RAID5s. In late 2013 we added a third storage server: a custom-built supermicro box that supports 45 disks and up to 150TB storage split into 3 RAID arrays. As of 2013-12, this system partially populated to provide 48TB of storage. At USC/HPCC as DH: we operate on USC's High Performance Computing Cluster (details at <http://www.usc.edu/hpcc>). Among supercomputers in an academic setting, HPCC's supercomputer cluster is the 10th fastest in the United States. Among all supercomputers in the world, it is ranked 166th. It claimed this distinction by achieving a benchmark in Fall 2012 of 149.9 teraflops, or 149.9 trillion floating-point calculations per second, on its 2,225-node, 10-gigabit backbone cluster. HPCC comprises a diverse mix of computing and data resources. Two Linux clusters constitute the principal computing resource, with 785 dual-core Intel Xeons and 1990 quad-core Opterons. In addition, HPCC has a central facility that provides more than 1.2 petabytes of combined disk storage and potential access to nearly a petabyte of tape storage. Of this facility, LACREND maintains 26TB of dedicated disk storage. These facilities are used for processing anonymized packet header data and are planned to be used for researcher-to-data remote access.
- At CSU as DH: we operate three capture points, a data archive and compute nodes to support researcher- to-data operations. All three capture points run at 10Gb/s and they continuously capture traffic data, anonymized it and archive it in flow form. We have been archiving data since 2008. The archive contains about 180TB of storage, with another approximately 40TB distributed over various machines. 90 out of the 180TB were added in 2014-Q1. The compute node is a machine with 32GB of RAM, 4TB of storage and an additional 30TB in a directly attached RAID box. We received 10 more machines from DHS and we are adding them on-demand. In addition, we added two Dell PE720xd compute servers each with 128GB of RAM and 48TB of storage. Currently we use these for internal projects, but have periodically given access to affiliated groups.
- At USC/ISI and USC/ HPCC as DP for anonymized packet headers: LACREND does not directly operate any packet collection facilities, but USC/Information Technology Services operates LACREND hardware that provides anonymized packet headers. These facilities include three Dell packet capture boxes with Endace capture cards operating at 1Gb/s and 10Gb/s.
- At USC/ISI, CSU, and Keio University as DP for internet address space data collection: LACREND operates Internet Address probing boxes. These boxes vary, but are at least Dell, 4-core Intel Xeon or AMD Opterons, typically with 2 or more network interfaces. We currently do address

probing from three sites: USC/ISI, CSU, and Keio University's Shonan Fujisawa Campus. We are planning to deploy new address probing boxes at additional locations over the course of LACREND's period of performance. We have been working for some time to deploy a fourth site in Europe and hope to complete such plans in 2014.

- At CSU and FRGP as DP for packet flows: LACREND operates three 10Gb/s capture points. These are equipped with high-end Endace cards that allow line speed capture and ns time precision. Custom LACREND software merges the two flow directions and anonymizes the traffic. Our custom software also manages a RAM disk for cases where line speed operations need to be performed.

Software

- All of these computers run CentOS, Fedora, or Ubuntu versions of Linux.
- Data collection (the DP role) at USC/ISI, USC/HPCC, CSU, and FRGP is done with the LANDER packet processing software, described in the paper: "Experiences with a Continuous Network Tracing Infrastructure" by Alefiya Hussain, Genevieve Bartlett, Yuri Pradkin, John Heidemann, Christos Papadopoulos, and Joseph Bannister, In Proceedings of the ACM SIGCOMM MineNet Workshop, pp. 185-190. Philadelphia, PA, USA, ACM. August, 2005.
<<http://dx.doi.org/10.1145/1080173.1080181>>.
- Data distribution (the DH role), is done at USC/ISI using custom web-based software to provide secure access for specific datasets to researchers, handling of dataset requests, and auditing of the process. In 2013 we shifted to using Trac (an open source issue tracking system, see <http://trac.edgewall.org>) to handle dataset processing. We have found it works well to handle the complexities of multiple dataset requests in many different states, each waiting on feedback from different parties.
- For some sensitive datasets, data "distribution" (the DH role), is done using researcher-to-data, where researchers are given accounts and allowed to log-in to LACREND-operated machines. This model is currently operational at CSU, and it is planned for use at USC/HPCC

Backups, mirroring, and redundancy

- We back up all archival data at least once.
- Backups at USC/ISI are done either: for rapidly changing information, to an external RAID operated by ISI operations staff, or (beginning in 2013) static datasets have been backed up to external USB disks.
- Our use of external USB disk was added as a low-cost method for off-line storage, following this success at CSU. Both master copies and backups are all checksummed and checksums for one dataset is checked each day.
- Backups at USC/HPCC are done by their operations staff using off-line tape storage.
- Backups at CSU are done via external USB drives as needed. The archive runs RAID5 for robustness, and the external drives are used to backup all the data. The drives are then stored separately from the archive.

USC-ISI 2014

1. Description

- To operate as a PREDICT Data Host (DH), the LACREND project maintains computers at USC/ISI, USC/HPCC and CSU. To operate as a PREDICT Data Provider (DP), the LACREND project operates computers at USC/ISI, USC/HPCC, CSU, FRGP, and Keio University.

2. System Inventory

Hardware

- At USC/ISI as DH: primary dataset distribution is done on a webserver (a Dell 4-core Intel Xeon box running Fedora 19 Linux). Front-end web service on ISI's commodity web server (a Solaris Sun-280R running Solaris 2.9). Storage is handled by two storage servers. Three are Dell 2950s with fully populated internal disks running as JBOD, and with external RAIDs. In late 2013 we added a third storage server: a custom-built supermicro box that supports 45 disks and up to 150TB storage split into 3 RAID arrays. As of 2013-12, this system partially populated to provide 48TB of storage. As of 2014-12 we have built out a second rank of the system, and we expect to completely populate it in 2015.
- At USC/HPCC as DH: we operate on USC's High Performance Computing Cluster (details at <http://www.usc.edu/hpcc>). Among supercomputers in an academic setting, HPCC's supercomputer cluster is the 10th fastest in the United States. Among all supercomputers in the world, it is ranked 166th. It claimed this distinction by achieving a benchmark in Fall 2012 of 149.9 teraflops, or 149.9 trillion floating-point calculations per second, on its 2,225-node, 10-gigabit backbone cluster. HPCC comprises a diverse mix of computing and data resources. Two Linux clusters constitute the principal computing resource, with 785 dual-core Intel Xeons and 1990 quad-core Opterons. In addition, HPCC has a central facility that provides more than 1.2 petabytes of combined disk storage and potential access to nearly a petabyte of tape storage. As of 2013, at this facility, LACREND maintains 26TB of dedicated disk storage. These facilities are used for processing anonymized packet header data and are planned to be used for researcher-to-data remote access. As of 2014 we have reduced RAID disk space to 10TB, and added an 8-node Hadoop cluster with 16 TB of raw disk space (thus TB usable storage after replication). At CSU as DH: we operate three capture points, a data archive and compute nodes to support researcher-to-data operations. All three capture points run at 10Gb/s and they continuously capture traffic data, anonymized it and archive it in flow form. We have been archiving data since 2008. The archive contains about 180TB of storage, with another approximately 40TB distributed over various machines. 90 out of the 180TB were added in 2014-Q1. The compute node is a machine with 32GB of RAM, 4TB of storage and an additional 60TB in two directly attached RAID boxes. We received 10 more machines from DHS and we are adding them on-demand. In addition, we have two Dell PE720xd compute servers each with 128GB of RAM and 48TB of storage. Currently we use these for internal projects, but have periodically given access to affiliated groups.
- At USC/ISI and USC/ HPCC as DP for anonymized packet headers: LACREND does not directly operate any packet collection facilities, but USC/Information Technology Services operates LACREND hardware that provides anonymized packet headers. These facilities include three Dell packet capture boxes with Endace capture cards operating at 1Gb/s and 10Gb/s.

- At USC/ISI, CSU, Keio University, and Athens University of Economics and Business as DP for internet address space data collection: LACREND operates Internet Address probing boxes. These boxes vary, but are at least Dell, 4-core Intel Xeon or AMD Opterons, typically with 2 or more network interfaces. We currently do address probing from three sites: USC/ISI, CSU, and Keio University's Shonan Fujisawa Campus. We are planning to deploy new address probing boxes at additional locations over the course of LACREND's period of performance. In 2014 we deployed a fourth site in Greece at the Athens University of Economics and Business.
- At CSU and FRGP as DP for packet flows: LACREND operates three 10Gb/s capture points. These are equipped with high-end Endace cards that allow line speed capture and ns time precision. Custom LACREND software merges the two flow directions and anonymizes the traffic. Our custom software also manages a RAM disk for cases where line speed operations need to be performed.

Software

- All of these computers run CentOS, Fedora, or Ubuntu versions of Linux.
- Data collection (the DP role) at USC/ISI, USC/HPCC, CSU, and FRGP is done with the LANDER packet processing software, described in the paper: "Experiences with a Continuous Network Tracing Infrastructure" by Alefiya Hussain, Genevieve Bartlett, Yuri Pradkin, John Heidemann, Christos Papadopoulos, and Joseph Bannister, In Proceedings of the ACM SIGCOMM MineNet Workshop, pp. 185-190. Philadelphia, PA, USA, ACM. August, 2005. <<http://dx.doi.org/10.1145/1080173.1080181>>. We continue to evolve this software, partially as joint work with the DHS-sponsored Retro-future project started in 2012. As of 2014 we have added the use of HDFS for temporary storage of collected data, and we expect to support HDFS for both initial and processing storage in coming years.
- Data distribution (the DH role), is done at USC/ISI using custom web-based software to provide secure access for specific datasets to researchers, handling of dataset requests, and auditing of the process. In 2013 we shifted to using Trac (an open source issue tracking system, see <http://trac.edgewall.org>) to handle dataset processing. We have found it works well to handle the complexities of multiple dataset requests in many different states, each waiting on feedback from different parties. We have regularly updated our use of Trac, in 2014 integrating it with our data distribution back-end software.
- For some sensitive datasets, data "distribution" (the DH role), is done using researcher-to-data, where researchers are given accounts and allowed to log-in to LACREND-operated machines. This model is currently operational at CSU, and it is planned for use at USC/HPCC.

Backups, mirroring, and redundancy

- We back up all archival data at least once.
- Backups at USC/ISI are done either: for rapidly changing information, to an external RAID operated by ISI operations staff, or (beginning in 2013) static datasets have been backed up to external USB disks.
- Our use of external USB disk was added as a low-cost method for off-line storage, following this success at CSU. Both master copies and backups are all checksummed and checksums for one dataset is checked each day. In 2014 USC/ISI switched from off-line backups with USB to on-line,

but off-site backups using a RAID deployed at CSU. We find this deployment much easier to use and not much more expensive than off-line USB-based backups.

- Backups at USC/HPCC are done by their operations staff using off-line tape storage.
- Backups at CSU are done via external USB drives as needed. The archive runs RAID5 for robustness, and the external drives are used to back up all the data. The drives are then stored separately from the archive.

USC-ISI 2015

1. Description

- To operate as a PREDICT Data Host (DH), the LACREND project maintains computers at USC/ISI, USC/HPCC and CSU. To operate as a PREDICT Data Provider (DP), the LACREND project operates computers at USC/ISI, USC/HPCC, CSU, FRGP, and Keio University.

2. System Inventory

Hardware

- At USC/ISI as DH: primary dataset distribution is done on a webserver (a virtual machine running Fedora 22 hosted on an 8-core Dell box running CentOS 7). Front-end web service is hosted on the same platform. Storage is handled by three storage servers. Three are Dell 2950s with fully populated internal disks running as JBOD, and with external RAID. We also support a custom-built Supermicro storage RAID that supports 45 disks and ~200TB storage split into 3 RAID arrays. As of 2015-12, this system is now fully populated and we have plans to purchase a second in 2016q1
- At USC/HPCC as DH: we operate on USC's High Performance Computing Cluster (details at <http://www.usc.edu/hpcc>). Among supercomputers in an academic setting, As of Nov. 2015, HPCC's supercomputer cluster is the 145th fastest in the world, based on 2015 benchmarks of 11,000 cores, reaching 555 TFLOP/s (Rmax). The current cluster has about 3400 cores, and the 11,000-core cluster is now the secondary cluster.
- LACREND purchases storage from the USC Digital Repository, operator of the Shoah Foundation a multi-media, digital repository of the Holocaust and other genocides. This storage is provided as a service and is tightly integrated with HPC.
- USC/ITS operates an 8-node Hadoop cluster, hosted at USC. As of 2015 this cluster is used primarily for storage. We are in the process of shifting it to use for both compute and storage.
- At USC/ISI and USC/ HPCC as DP for anonymized packet headers: LACREND does not directly operate any packet collection facilities, but USC/Information Technology Services operates LACREND hardware that provides anonymized packet headers. These facilities include three Dell packet capture boxes with Endace capture cards operating at 1Gb/s and 10Gb/s.
- Capture service was suspended mid-2015 due to USC's upgrade of the campus Internet connection to a single 100 Gb/s link. LACREND cannot afford capture hardware that will sustain this bitrate, and even if we could, USC does not have sufficient ports to provided data. Instead, we began plans to transition to capture a 10Gb/s link internal to USC that connects its wireless networks to the campus. This work was approved and begun March 2015 and hardware was provided to ITS in June 2015. As of Dec. 2015 this hardware has been placed in-line, but capture is not yet active and integration continues.
- At USC/ISI, CSU, Keio University, and Athens University of Economics and Business as DP for internet address space data collection: LACREND operates Internet Address probing boxes. These boxes vary, but are at least Dell, 4-core Intel Xeon or AMD Opterons, typically with 2 or more network interfaces. We currently do address probing from three sites: USC/ISI, CSU, and Keio University's Shonan Fujisawa Campus. We are planning to deploy new address probing boxes at additional locations over the course of LACREND's period of performance. In 2014 we deployed a fourth site in Greece at the Athens University of Economics and Business.

- In 2015 we began plans to upgrade the computer hardware at Keio University, replacing the single 2U capture machine with two faster, 1U machines.
- At CSU and FRGP as DP for packet flows: LACREND operates three 10Gb/s capture points. These are equipped with high-end Endace cards that allow line speed capture and ns time precision. Custom LACREND software merges the two flow directions and anonymizes the traffic. Our custom software also manages a RAM disk for cases where line speed operations need to be performed. In 2015 we added a Netflow feed from FRGP that captures a sampled picture from multiple routers. We wrote custom software to anonymize the Netflow feed using prefix-reserving anonymization. We plan to make this data available along with our other flow data as soon as we are happy with the stability of the feed.

Software

- All of these computers run CentOS, Fedora, or Ubuntu versions of Linux.
- Data collection (the DP role) at USC/ISI, USC/HPCC, CSU, and FRGP is done with the LANDER packet processing software, described in the paper: "Experiences with a Continuous Network Tracing Infrastructure" by Alefiya Hussain, Genevieve Bartlett, Yuri Pradkin, John Heidemann, Christos Papadopoulos, and Joseph Bannister, In Proceedings of the ACM SIGCOMM MineNet Workshop, pp. 185-190. Philadelphia, PA, USA, ACM. August, 2005. <<http://dx.doi.org/10.1145/1080173.1080181>>. We continue to evolve this software, partially as joint work with the DHS-sponsored Retro-future project started in 2012.
- As of 2014 we have added the use of HDFS for temporary storage of collected data, and we expect to support HDFS for both initial and processing storage in coming years.
- In 2015 we took steps to improve our use of Hadoop with "hard-link emulation". This step will allow use of Hadoop storage at all levels of LANDER packet processing, rather than just as the initial data repository. Testing of this software started in 2015q4 and we expect operational deployment in 2016.
- Data distribution (the DH role), is done at USC/ISI using custom web-based software to provide secure access for specific datasets to researchers, handling of dataset requests, and auditing of the process. In 2013 we shifted to using Trac (an open source issue tracking system, see <http://trac.edgewall.org>) to handle dataset processing. We have found it works well to handle the complexities of multiple dataset requests in many different states, each waiting on feedback from different parties. We have regularly updated our use of Trac, in 2014 integrating it with our data distribution back-end software.
- For some sensitive datasets, data "distribution" (the DH role), is done using researcher-to-data, where researchers are given accounts and allowed to log-in to LACREND-operated machines. This model is currently operational at CSU, and it is planned for use at USC/HPCC.

Backups, mirroring, and redundancy

- We back up all archival data at least once.
- Backups at USC/ISI are done either: for rapidly changing information, to an external RAID operated by ISI operations staff, or (beginning in 2013) static datasets have been backed up to external USB disks.
- Our use of external USB disk was added as a low-cost method for off-line storage, following this success at CSU. Both master copies and backups are all checksummed and checksums for one

dataset is checked each day. In 2014 USC/ISI switched from off-line backups with USB to on-line, but off-site backups using a RAID deployed at CSU. We find this deployment much easier to use and not much more expensive than off-line USB-based backups.

- In 2015q4 we started building out the ISI backup RAID hosted at CSU to full capacity, with plans to add another 95TB of storage as of 2016q1.
- Backups at USC/HPCC are done by their operations staff using off-line tape storage.
- Backups at CSU are done via external USB drives as needed. The archive runs RAID5 for robustness, and the external drives are used to backup all the data. The drives are then stored separately from the archive. In the past year we switched to RAID with 8TB hard drives. We were able to reuse our RAID boxes and quadruple their capacity with the new disks.

Storage Summary	Storage Type	Used (TB)	Available (TB)	Capacity (TB)	Partitions
The following table summarizes project storage as of Dec. 2015.Site					
isi	hdfs	82.4	20.1	102.5	1
isi	jbod	35.5	8.6	44.1	14
isi	raid	196.56	44.7	241.3	6
isi	subtotal	258	227	486	21
USC	raid	938.6	333.9	1275	1
USC	Hdfs	82.7	16.1	98.9	1
CSU	archive	-	-	180	-
CSU	External researchers			52TB	6
CSU	Capture boxes			40TB	4
LACREND	total	-	-	726TB	-