

PCH Hosting Infrastructure Reports

PCH 2012

1. Description

- Packet Clearing House's (PCH) PREDICT datasets are maintained on a centralized storage and database platform. We maintain three Promise RAID storage servers hosted in two interconnected datacenters in Sunnyvale and Palo Alto. Combined, these RAID servers give PCH a capacity of 70TB of available storage capacity.

2. System Inventory

- Ongoing data-collection of our Routing Topology, Prefix origin inconsistency, and MRT Datasets is controlled by a server that mounts the primary storage RAID over NFS for read and write access. Routing topology data from our operational nodes is retrieved daily and archived by this server, along with the synthetic datasets produced by DARPA and IARPA. This data is made available to authenticated through a web-server which also mounts the storage over NFS.
- The Internet infrastructure topology and DNS infrastructure datasets produced by PCH are stored in our database systems. These are primarily stored on a database server running MySQL. Application specific data from the INOC-DBA system is stored on the host server for that application and replicated to our main database server. Domain Name System (DNS) query metadata are collected and stored on a separate MySQL server that is dedicated specifically to this resource-intensive task.

PCH 2013

1. Description

- Packet Clearing House's (PCH) PREDICT datasets are maintained on a centralized storage and database platform. We maintain three Promise RAID storage servers hosted in two interconnected datacenters in Sunnyvale and Palo Alto. Combined, these RAID servers give PCH a capacity of 70TB of available storage capacity.

2. System Inventory

- Ongoing data-collection of our Routing Topology, Prefix origin inconsistency, and MRT Datasets is controlled by a server that mounts the primary storage RAID over NFS for read and write access. Routing topology data from our operational nodes is retrieved daily and archived by this server, along with the synthetic datasets produced by DARPA and IARPA. This data is made available to authenticated through a web-server which also mounts the storage over NFS.
- The Internet infrastructure topology and DNS infrastructure datasets produced by PCH are stored in our database systems. These are primarily stored on a database server running MySQL. Application specific data from the INOC-DBA system is stored on the host server for that application and replicated to our main database server. Domain Name System (DNS) query metadata are collected and stored on a separate MySQL server that is dedicated specifically to this resource-intensive task.
- Because our PREDICT funding was decreased, we have used overhead and complementary funding sources to cover the cost of a new distributed file system running on top of four Cisco UCS C240 M3 LFF servers with, collectively, 144TB of disk, 160 cores of CPU, and 1TB of RAM.

PCH 2014

1. Description

- Packet Clearing House's (PCH) PREDICT datasets are maintained on a centralized storage and database platform. We maintain three Promise RAID storage servers hosted in two interconnected datacenters in Sunnyvale and Palo Alto. Combined, these RAID servers give PCH a capacity of 70TB of available storage capacity.

2. System Inventory

- Ongoing data-collection of our Routing Topology, Prefix origin inconsistency, and MRT Datasets is controlled by a server that mounts the primary storage RAID over NFS for read and write access. Routing topology data from our operational nodes is retrieved daily and archived by this server, along with the synthetic datasets produced by DARPA and IARPA. This data is made available to authenticated through a web-server which also mounts the storage over NFS.
- The Internet infrastructure topology and DNS infrastructure datasets produced by PCH are stored in our database systems. These are primarily stored on a database server running MySQL. Application specific data from the INOC-DBA system is stored on the host server for that application and replicated to our main database server. Domain Name System (DNS) query metadata are collected and stored on a separate MySQL server that is dedicated specifically to this resource-intensive task.
- A new distributed file system running on top of four Cisco UCS C240 M3 LFF servers with, collectively, 144TB of disk, 160 cores of CPU, and 1TB of RAM, noted as planned in our 2013 report, remains pending.

PCH 2015

1. Description

- Packet Clearing House's (PCH) PREDICT datasets are maintained on a centralized storage and database platform. We maintain three Promise RAID storage servers hosted in two interconnected datacenters in Sunnyvale and Palo Alto. Combined, these RAID servers give PCH a capacity of 70TB of available storage capacity.

2. System Inventory

- Ongoing data-collection of our Routing Topology, Prefix origin inconsistency, and MRT Datasets is controlled by a server that mounts the primary storage RAID over NFS for read and write access. Routing topology data from our operational nodes is retrieved daily and archived by this server, along with the synthetic datasets produced by DARPA and IARPA. This data is made available to authenticated through a web-server which also mounts the storage over NFS.
- The Internet infrastructure topology and DNS infrastructure datasets produced by PCH are stored in our database systems. These are primarily stored on a database server running MySQL. Application specific data from the INOC-DBA system is stored on the host server for that application and replicated to our main database server. Domain Name System (DNS) query metadata are collected and stored on a separate MySQL server that is dedicated specifically to this resource-intensive task.
- A new distributed file system running on top of four Cisco UCS C240 M3 LFF servers with, collectively, 144TB of disk, 160 cores of CPU, and 1TB of RAM, noted as planned in our 2013 report, remains pending.