

# CPR Status and Roadmap (August 2005).

The Campus-wide network Performance and Recovery (CPR) project [1] aims to provide quantitative data to resolve network problems. This document reviews the current status and future goals of the project.

## 1. Background

Troubleshooting network issues is hard. Typically, network engineers rely on years of experience to investigate problems. Tools such as ping and traceroute are being increasingly blocked in the name of security. The CPR project aims to track performance and provide alerts. It is unlikely that human experience can be replaced with an algorithm, but quantifying the problem and identifying its extent will improve the time taken to resolve many issues.

## 2. Status

The project has completed deployment phase. Management, analysis and visualization tools are being installed.

### 2.1. Current Deployment

53 hosts are currently deployed (or are in the process of being deployed). The bulk of the hardware deployed to date has been surplus. Consequently the low direct cost has been very low. Figure 1 shows the logical distribution across the Atlanta campus and beyond. In addition, cpr-homeland and cpr-central store and analyze the data.

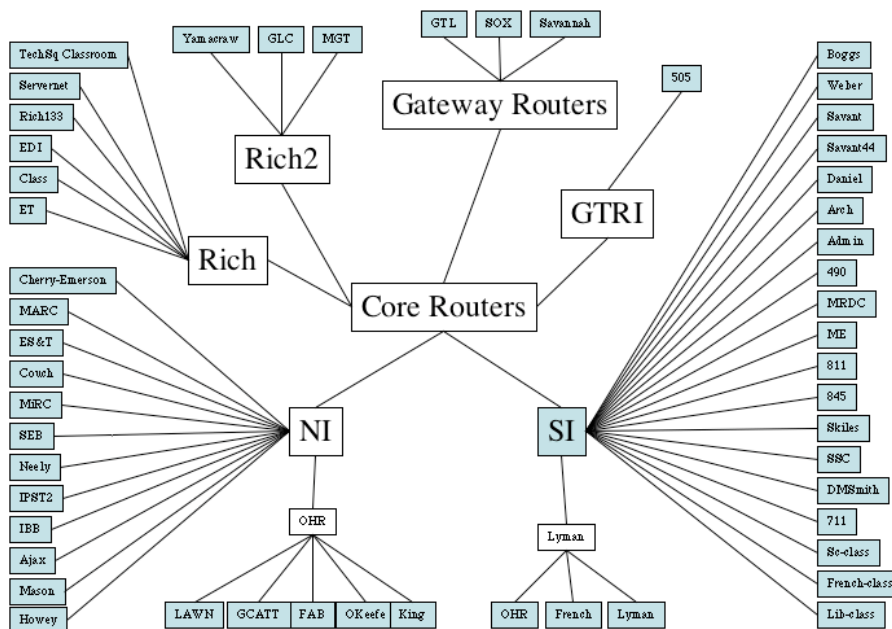


Figure 1: Current CPR deployment. The shaded boxes indicate the logical location where the monitors are deployed.

## 2.2. Case Studies

CPR has quickly become a well-used tool by the members of the network group responsible for trouble-shooting and fixing problems.

- Users reported they could not exchange email with collaborators at another University. Another group had investigated the issue before giving up and passing the ticket to the network group. A CPR host was already deployed on the subnet. Typically the network group would not have direct access to the subnet and this scenario would lead to finger pointing, particularly at the firewall. However, using the CPR node, the issue was pinpointed within 20 minutes. An ACL in another network between the Universities was blocking traffic from the specific subnet where the mail server was connected. However, it took several weeks to identify the right contacts at the network to have the ACL removed.
- Users in a particularly building reported losing connectivity to their home directories housed on a server in another building. A CPR host was deployed. Immediately the CPR host provided evidence that the server had been compromised and was running unauthorized software. It was removed from service, cleaned and returned to service. The disconnections ceased. There are sporadic reports of users being disconnected from the network. In another case users report being disconnected from sessions within the same subnet. CPR has been used to quantify the issue by tracking the hosts being disconnected and looking for patterns, for example if the disconnected hosts are all on the same switch. However, in this case no pattern has been seen. This issue remains unresolved. It provides motivation to extend CPR to include layer 2 tools. Other cases have been resolved when CPR indicated the problem was restricted to one particular switch, or only DHCP clients were affected. Identifying what to reboot is a significant contribution to the work of the network team.
- Users in several buildings reported problems accessing the central email servers. CPR collects data on the TCP round trip time and the application connection time. In this case no disruption was observed indicating that the issue was deeper in the client-server exchange than CPR probed.
- One particular user complained about poor throughput between two buildings. CPR hosts were deployed on both subnets. CPR confirmed poor throughput, eliminating the possibility of mis-configuration on the users host. Coincidentally, errors were observed on the switchport leading to the diagnosis of a duplex mismatch.
- Georgia Tech's satellite campus in Savannah experienced loss of connectivity. The CPR host in Savannah tracks the routing and could have identified the exact location and nature of the problem except it was unreachable. Furthermore, for security reasons, traceroute is blocked at the entrance to Savannah so it is not possible to determine the problem from the Georgia Tech side. However, the

traceroute does make it to the edge of the peachnet network, indicating the problem might have been within the Savannah campus network.

### 2.3. Measurements

The smokeping [2] and nagios [3] tools are installed on all CPR hosts. Smokeping measures and creates graphs of the round trip delay between hosts. Nagios tracks the response of applications such as the central mail servers.




Figure 2: Smokeping graph showing round trip time between two CPR hosts.

Traceroute is also run between all CPR hosts. Within campus, it is unlikely that the route will change. However, tracking it provides a sanity check and should be an early indication of serious problems on the backbone if anything does appear to change.

The Pathload tool [4] has been tested on some CPR hosts. It provides information on the available bandwidth. It may be a real-time, low-impact indication of high bandwidth utilization, such as with a denial of service attack.

The Internet2 e2epi tools; bwctl, owamp and ndt [5] are also installed on some hosts. However, because of problems with the accuracy of NTP and technical issues with the patched kernel, testing and evaluation has not taken place.

### 2.4. Impact of Monitoring

The CPR hosts can easily drive 100 Mbps of network throughput across campus. However, the CPU usage required for this impacts ping measurements. This impact looks like high jitter on the network, which does not exist. The Internet2 end-to-end group experienced similar problems running their bwctl and owamp tools on the same hardware. Their solution is to deploy the tools on separate hardware. This is not really an option for CPR. Instead we will develop a meta-scheduler to make sure the tests run at different times.

### 2.5. Security

From the initial concept, it was planned that CPR would be a platform for other tools and projects. The security tools Nessus and nmap are installed on all CPR hosts. The information security team can monitor activity within the subnet, behind the firewall. Also their monitoring can be distributed and not have to run for very long periods of time, impacting the network.

### 2.6. Management

As the number of CPR hosts increase, the challenge of maintaining them becomes harder. Each time a new host is added, or even a minor configuration change has to be made, all

the hosts have to be touched. Often several files will be involved. The solution is to create a tool capable of distributing changes from a central server to all the deploy hosts. Existing tools developed within the college of computing were considered. However, their tools use network connectivity. Part of CPR is to allow for network problems. In addition, the maintenance of the infrastructure is a research project in itself, and it was decided to experiment with the emerging standards and develop a service oriented architecture.

Furthermore, because the hosts are surplus equipment, it was decided they should be considered to be disposable. Reliability has been an issue. The information security team has set up a separate nagios server to monitor the CPR hosts. Figure 3 shows part of a table from the nagios report. It is expected that the CPR project will demonstrate tremendous value and new equipment can be budgeted for. Obviously this equipment will not be considered disposable.

<a href="#">cpr-servernet</a>	UP	5 OK	
<a href="#">cpr-si</a>	UP	5 OK	
<a href="#">cpr-skiles</a>	UP	5 OK	
<a href="#">cpr-sox</a>	UP	4 OK 1 CRITICAL	
<a href="#">cpr-ssc</a>	UP	5 OK	
<a href="#">cpr-techsqclass</a>	UP	5 OK	
<a href="#">cpr-weber</a>	UP	5 OK	
<a href="#">cpr-yamacraw</a>	UP	5 OK	

**Figure 3: Part of the Nagios table showing processes running on CPR hosts. 5 services are monitored; smoking, nagios, ntp, ssh and http. This screenshot was taken when one of these services was unresponsive on cpr-sox.**

## 2.7. Access Policy

CPR is primarily an active-measurement platform. Sensitive data involving user activity is not gathered. Access to certain results should be openly available. One of the design goals was to make CPR useful to the CSRs and the end users, as well as the network support team.

Login (ssh) access to the machines is restricted to the development team, network support, information security and operations. Access to other groups such as facilities has yet to be defined.

Currently, hosts connected to the same subnet as a CPR host can view the measurements directly from the local web server. General access to graphs is available on the central server.

## 2.8. Data Backup

Data is frequently uploaded to the central servers using scp. Backup to tape or DVD will be evaluated.

### **3. Roadmap**

CPR is designed to be a production quality tool for trouble-shooting. It is also a platform for research and other projects.

#### **3.1. Further Deployment**

At the time of writing, the campus network consists of 164 network buildings on campus. There are 806 stacks of switches and 52440 ports. The team aim to deploy a host in every building, and perhaps located with every stack. The CPR team will continue to deploy donated machines in closets around campus. Purpose-bought machines would accelerate the deployment and provide increased stability and reliability. The team has investigated the use of Linksys wireless routers and macmini computers as CPR devices. CPR devices should be standard with new network equipment. The logistics of deploying CPR hosts in Singapore and Shanghai is being investigated.

#### **3.2. Development**

Deployment and maintenance are critical aspects of the CPR platform but software development is the major goal for CPR. However, development of layer 3 measurement tools is not part of the CPR project. The team will incorporate existing tools in the CPR framework.

- The current measurements may be supplemented. Researchers have developed numerous tools to measure aspects of network performance. Assessing these tools and their ability to help troubleshoot and/or diagnose problems would be a valuable task. Furthermore, these tools are typically run on test-beds or small deployments across the Internet. Deploying them on the CPR infrastructure would provide valuable feedback to the developers.
- The current measurements may be changed. Quantifying basic layer3 connectivity is a core goal for CPR. There are numerous tools available. These tools can be assessed and the one most useful for campus needs will be selected. Smokeping was chosen because it includes a graphical output. With additional visualization tools, it may be possible to use fping or ping instead. In addition to tools to measure round trip times, there are also tool to measure on-way delay.
- Problems are not restricted to Layer 3. Especially on campus, layer 2 measurements would be particularly useful. For example, tools exist that would allow the CPR team to track spanning tree maps. The netdisco tool has been used in other projects. Incorporating this may be useful.
- Monitoring services with application-layer tools such as Nagios has provided useful trouble-shooting information. Additional services will be monitored, for example DHCP.
- Analysis of the data is a central goal for CPR. Viewing graphs of performance is not scalable. Smokeping and Nagios send email to notify the administrator of

problems. However, a flood of email may indicate a server had been rebooted or a major router has crashed. Analysis will correlate measurements and indicate the nature and location of the issue. The network research community has developed statistical models of Internet performance. In addition to trouble shooting tools, CPR can identify statistical descriptions of base line performance.

- Network administrators have developed other tools to help trouble-shoot problems. One of the original goals for CPR was to utilize information gathered by the tools used at Georgia Tech. The GOAT tool in particular. The team may also investigate the feasibility of collecting netflow data for the local subnet. It would also be useful to gather information using SNMP and look for interface errors that would impact performance. It might also be useful to periodically scan logs to look for vlan or duplex mismatches.
- Development of inter-CPR communication is also a critical development project. Initial development has been conducted using the SOAP protocol. A complete service oriented architecture will be designed. The principal goal is to interoperate with the emerging International measurement infrastructure using standards defined by the Global Grid Forum (GGF) Network Monitoring Working Group (NMWG). This approach will allow CPR to probe networks connecting the Georgia Tech International Campuses. Development of a SOAP-based communication system will also be useful for 3<sup>rd</sup> party applications.
- Visualization is also an important aspect of the project. A sophisticated graphing tool has been developed and will be released to beta testers. Other visualization tools will be developed. The NLANR AMP project has a large suite of visualization tools. The team will deploy and assess the AMP package.
- The CPR team will develop a tool similar to the LAWN weedeater tool. It will assist end users in trouble shooting problems. The user will be able to visit a web page that will analyze the performance to their host and identify any issues. If an issue is found, a report can be sent to the appropriate group.
- CPR may also be a platform used to support and monitor performance on the wireless network.

It is hard to estimate the timeline for the development of the various parts of CPR. Use of students drives the development. CPR projects is also a useful experience for the students.

### **3.3. Other Projects**

The CPR project can easily be extended to IPv6. The Georgia Measurement and Monitoring (GAMMON) projects aims to extend the CPR concept throughout the state.

Cingular have expressed an interest in working on monitoring cells sites.

#### **4. References**

- [1] The original CPR white paper is on line at <http://www.rnoc.gatech.edu/cpr/cpr.doc>
- [2] The Smokeping web site is <http://people.ee.ethz.ch/~oetiker/webtools/smokeping/>
- [3] The Nagios web site is <http://www.nagios.org/>
- [4] The Pathload web site is <http://www.pathrate.org>
- [5] The Internet2 end-to-end performance initiative web site is <http://e2epi.internet2.edu>