

Title: TSUNAMI**Contacts:** Mark Meiss (mmeiss@indiana.edu)**URLs / RFCs / Papers**

[SUNAMI-03] Mark R. Meiss “TSUNAMI: A High-Speed Rate-Controlled Protocol for File transfer
[SUN-02] README file of tsunami -2002-12-02 release. <http://www.indiana.edu/~anml/anmlresearch.html>

Principle / Description of Operation

Current implementation of Tsunami contains two user-space applications, a client and a server. Tsunami uses UDP for data transfer and TCP for transferring control information such as retransmission request, restart request, error report, completion report. The overall architecture of the Tsunami is shown in the figure 1.

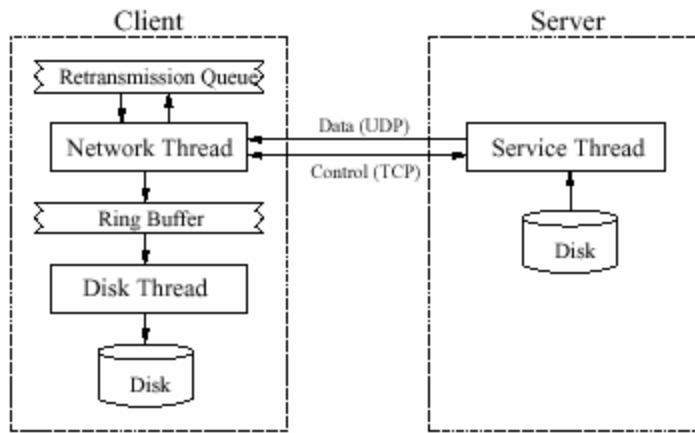


Figure 1: Tsunami architecture

In the initial negotiation phase, client sends its desired block size, target transfer rate, error threshold and inter-packet delay scaling factors. Server will respond the client with length of the file, agreed block size, the number of blocks and a time stamp. During the file transfer client uses two threads. Network thread to handle all network communication, maintain retransmission queue and to place blocks to ring buffer from disk storage. Disk thread moves the blocks from ring buffer to the target destination.

At the server side, server first processes the messages on TCP control channel; if there are no messages on the TCP control channel then server sends the next block.

At the client side, disk thread pulls the block out of the ring buffer. If the block has not been saved already, it is saved to disk and the bit array is updated to reflect the reception of the block and decrements the number of blocks remaining. More functionality and algorithm about server and client can be found at [SUNAMI-03]. The main advantage of Tsunami being, it allows user to tune the protocol.

Supported operation mode:

disk-to-disk (i.e. file transfer protocol, not general transport)

Authentication: Very simple authentication mechanism.

Implementations / API: No

Congestion Control Algorithms:

Limited. If the packet loss above the threshold causes exponential rise in inter packet delay and loss below the threshold causes exponential decrease in inter packet delay.

Fairness: Yet to be determined.

TCP Friendly: No

Predictable Performance Model: No

Results:

TSUNAMI recorded mean transmission rate of 850Mb/s for over 17 hours on the Global Terabit research network between Seattle, Washington and Brussels, Belgium.

Dedicated connection between TRIUMF in Vancouver, British Columbia, and CERN in Geneva, Switzerland with 150ms latency recorded the bandwidth of 600Mbps and 1Gbps achieved on occasion.

Target Usage Scenario:

Useful for bulk data transfer. Experiments have to be performed on high delay networks.