

Performance Evaluation of iSCSI-APT (iSCSI with Automatic Parallelism Tuning) on SINET3 with Layer-1 Bandwidth on Demand Service

* Hiroyuki Ohsaki, † Jakkapan Kittananun †, * Fumito Inoue, ‡ Yoshihiro Nomoto, * Makoto Imase

* School of Information Technology, King Mongkut's University of Technology, Bangkok, Thailand

† Graduate School of Information Science and Technology, Osaka University, Osaka, Japan

‡ NTT Service Integration Laboratories, NTT Corporation, Tokyo, Japan

Extended Abstract

In our previous work, we have proposed a parallelism tuning mechanism for the iSCSI protocol called *iSCSI-APT (iSCSI with Automatic Parallelism Tuning)*, which automatically adjusts the number of multiple connections according to network status. iSCSI-APT is primarily designed for bulk data transfer applications such as remote backup. iSCSI-APT measures the network status in a passive way, and automatically adjusts the number of multiple connections so that the iSCSI throughput is maximized.

In this paper, we evaluate the performance of iSCSI-APT on a *real* long-fat network, SINET3 (Science Information Network 3).

SINET3 is a nation-wide science information network launched in April 2007. SINET3 provides a variety of advanced network services since it has been designed primary for scientific researchers. One of the notable features of SINET3 for networking researchers is its *layer-1 bandwidth on demand service*. The layer-1 bandwidth on demand service allows SINET3 users to dynamically reserve/release a layer-1 path between SINET3 sites. We believe such an advanced network service in SINET3, the layer-1 bandwidth on demand service, is quite beneficial to networking researches because of its *reality*, *flexibility*, and *reproducibility*.

We perform a number of experiments on SINET3 for evaluating the performance of the iSCSI-APT while changing the bandwidth allocation between sites.

Figure 1 shows the network configuration used in our experiments. The iSCSI initiator is placed at Osaka University (Suita, Osaka). iSCSI targets are placed at NTT Service Integration Laboratories (Musashino, Tokyo) and Hokkaido University (Sapporo, Hokkaido). The iSCSI initiator at Osaka University and the iSCSI target at Hokkaido University are directly connected to SINET3 routers. The iSCSI target at NTT Service Integration Laboratories is connected to a SINET router at NII (National Institute of Informatics) (Chiyoda-ku, Tokyo) using a dedicated 1 [Gbit/s] link.

In all experiments, 10 [Gbyte] of data is transferred from the iSCSI initiator at Osaka University to either of iSCSI targets at NTT Service Initiator Laboratories and Hokkaido University. Note that the average round-trip time was approximately 12 [ms] between Osaka University and NTT Service Initiator Laboratories, and 37 [ms] between Osaka University and Hokkaido University.

We varied the bandwidth of the layer-1 path between Osaka University and NII, and between Osaka University and Hokkaido University using SINET3's layer-1 bandwidth on demand service. The bandwidth of the layer-1 path can be configured from 150 to 1,050 [Mbit/s] with 150 [Mbit/s] granularity.

Figure 2 compares iSCSI goodputs with and without iSCSI-APT (with a single connection) for different layer-1 path bandwidths. This figure shows that iSCSI-APT utilizes the network resource quite effectively regardless of the bottleneck bandwidth. On the contrary, the performance of iSCSI without iSCSI-APT is rather low, in particular, when the layer-1 path bandwidth is large.

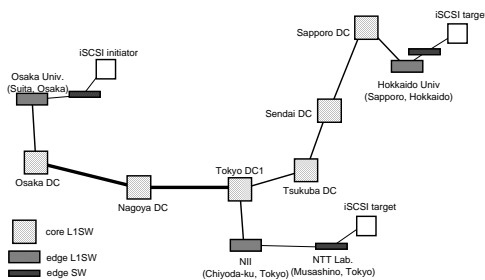


Figure 1: Network configuration used in experiments; the iSCSI initiator at Osaka University (Suita, Osaka) is connected to iSCSI targets at NTT Service Integration Laboratories (Musashino, Tokyo) and Hokkaido University (Sapporo, Hokkaido).

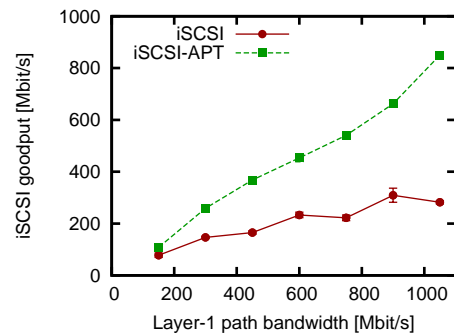


Figure 2: Layer-1 path bandwidth (Osaka University–NII, 150–1,050 [Mbit/s]) vs. iSCSI goodput; iSCSI-APT utilizes the network resource effectively regardless of the bottleneck bandwidth.