

Preliminary results of QoS performance measurement at CERN

March 2003

Japanese group

Contents

- Setup
- Event builder performance without QoS on 1x1 system (EoE rate and Throughput)
- Event builder performance without QoS on 6x1 system (EoE rate and Throughput)
- Event builder performance **with QoS** on 6x1 system (EoE rate and Throughput)

Setup

- pcatb137-146 with gigabit Ethernet NICs and a Switch (BATM)
- RedHat 7.2.1 with gcc-alt-2.95.2(CERN standard)
- Linux kernel2.4.9-31 with HZ4096 parameter
- DC software : DC-00-02-02
- Online Software : Online-00-17-02
- Kernel buffer size : 8MB

DC Parameters

ROSE

Number of RoB : 1

robDataSize : variable

SFI

MaxNumberOfAssembledEventsHigh : 50

MaxNumberOfAssembledEventsLow : 45

DefaultCredits : 10

TimeoutSetup : 500

TrafficShapingParameter : 30

DFM

Credits : 3

Timeout : 5000

MaxAssignQueueSize : 1000

ClearGroupSize : 50

DecisionGroupSize : 10

LVL2 decision : a LVL2 accept in a decision

Configuration

1x1:

A RootController, a DCcontroller, DFMAutotester7
and a L2SV on pcatb137

A DFM on pcatb138

A ROS on pcatb139

A SFI on pcatb146

6x1:

A RootController, a DCcontroller, DFMAutotester7
and a L2SV on pcatb137

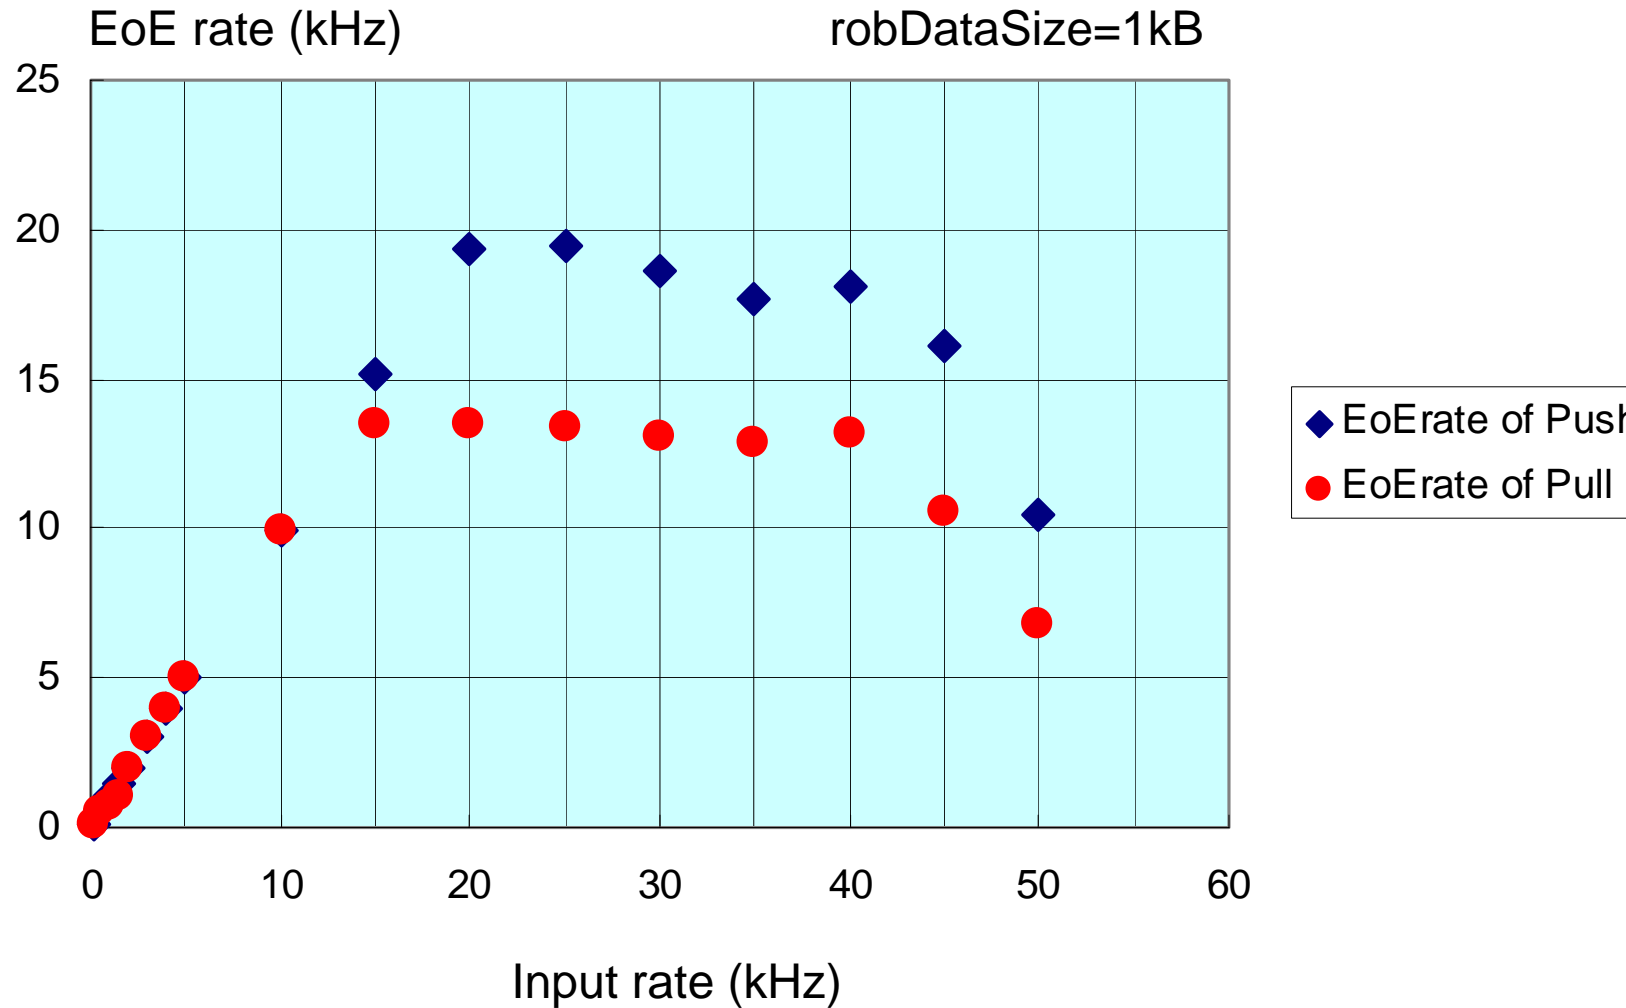
A DFM on pcatb138

6 ROSs on pcatb139-144

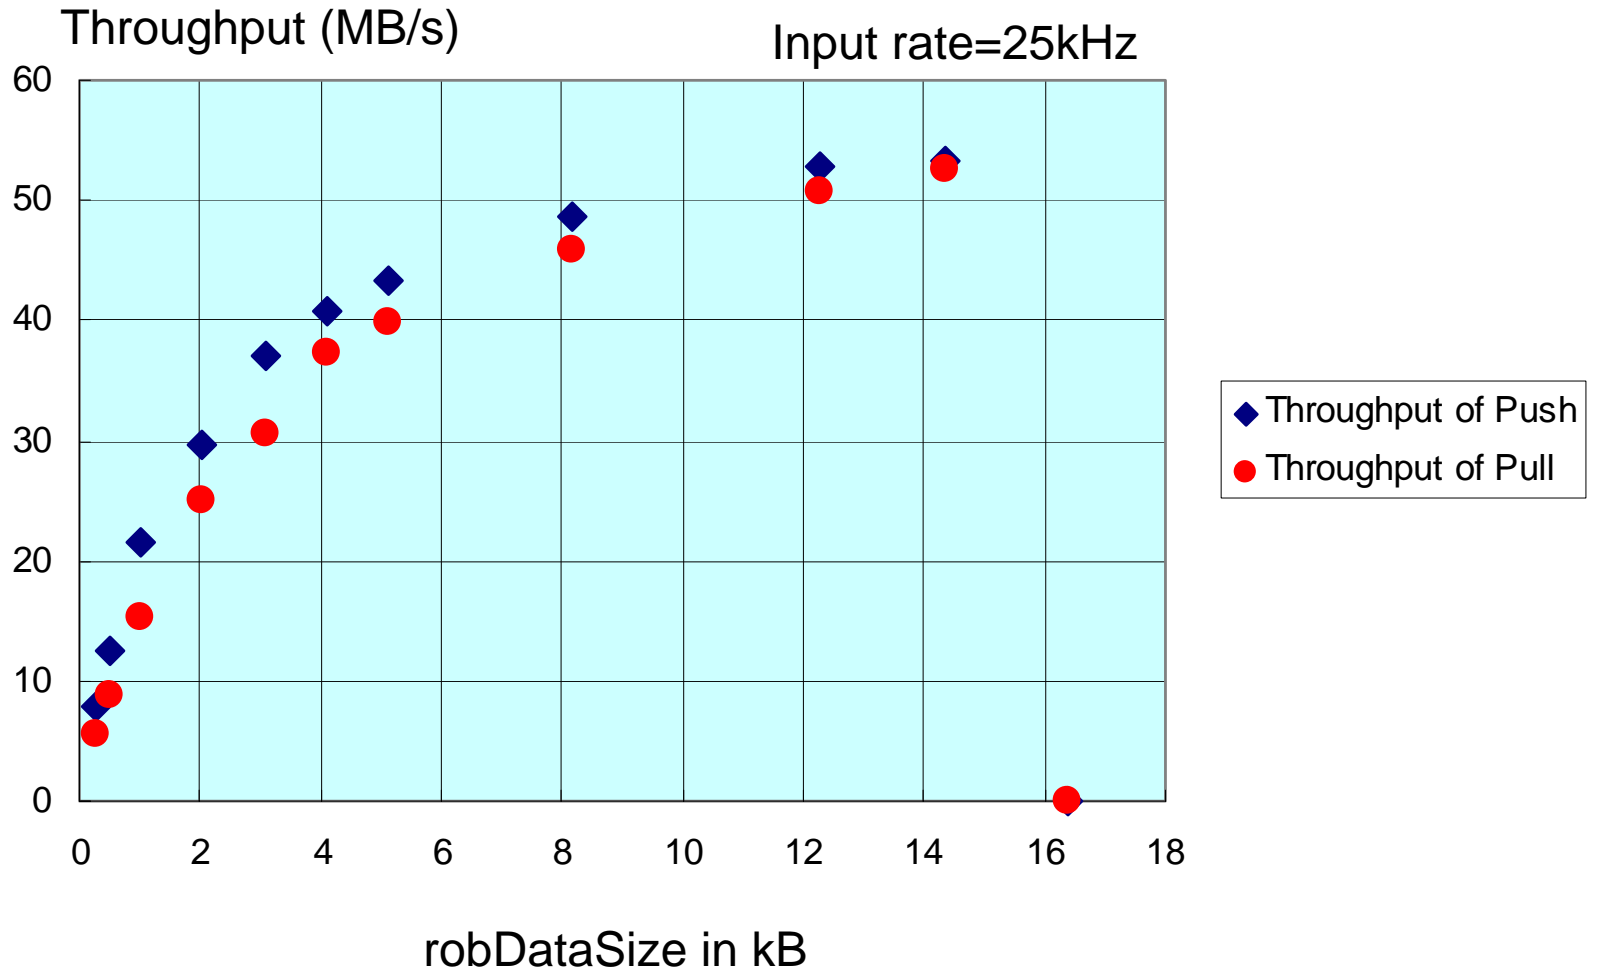
A SFI on pcatb146

1x1 system

EoE rate without QoS



Throughput without QoS




Results from the 1x1 system

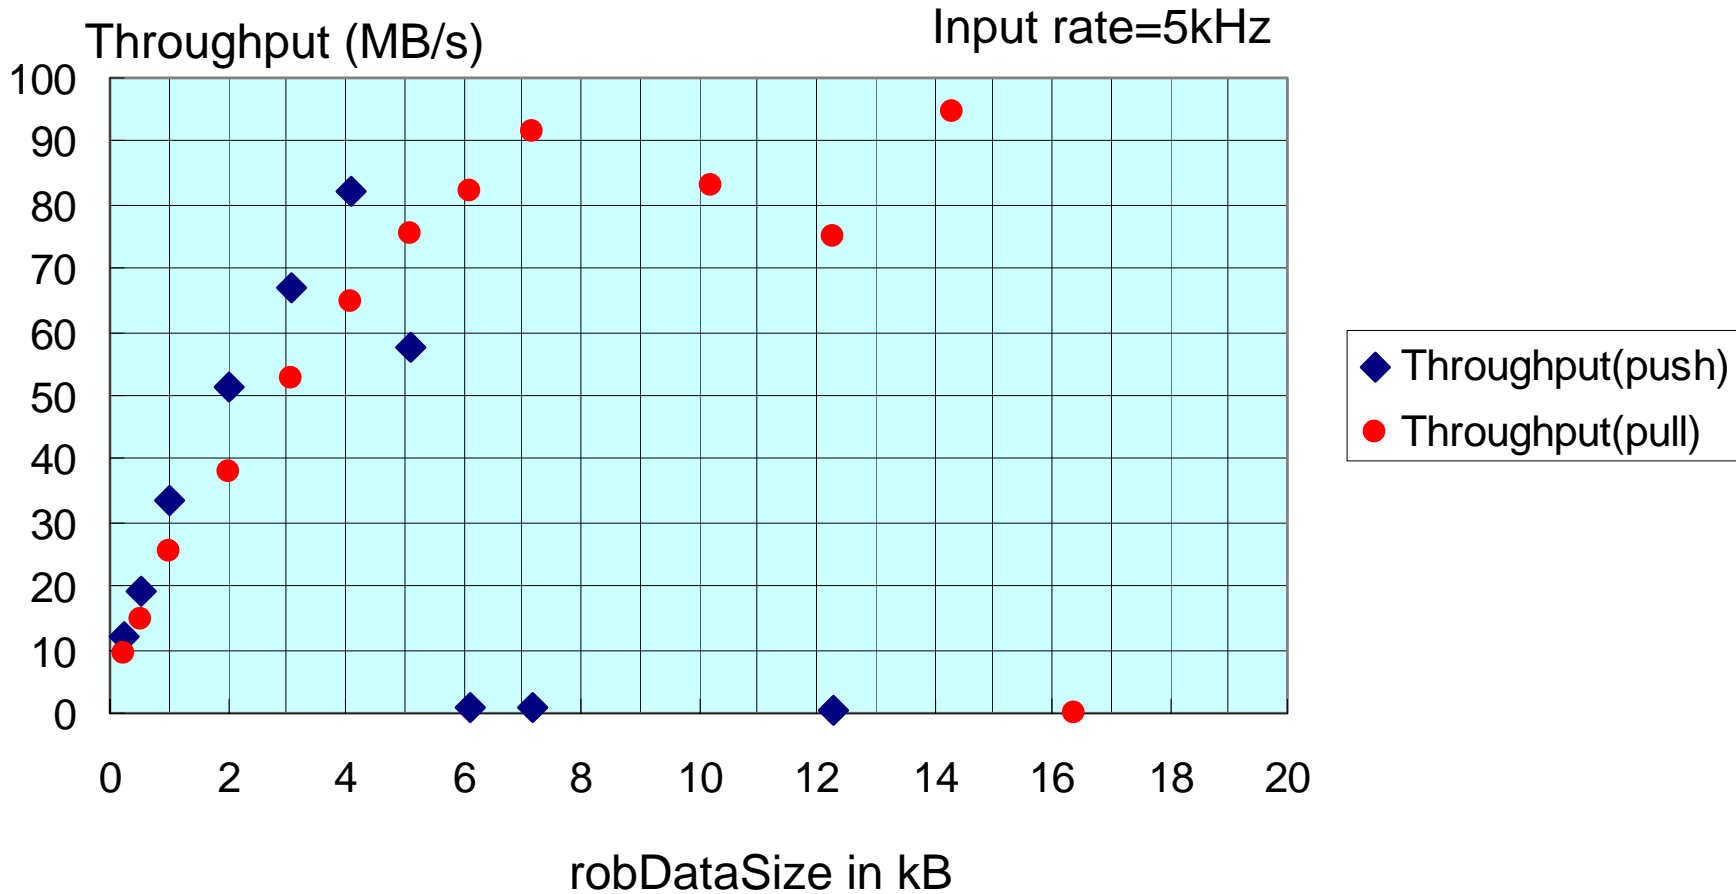
1. Performed push and pull scenarios with similar setups.
2. Push better than pull.
3. NO re-asks observed.
4. EoE rates dropped quickly with input trigger rates above 40-50 kHz.
5. Throughput reached over 50MB/s in push and pull scenarios, but the throughput dropped to zero at robDataSize of 16kB in both scenarios.

6x1 system

EoE rate without QoS



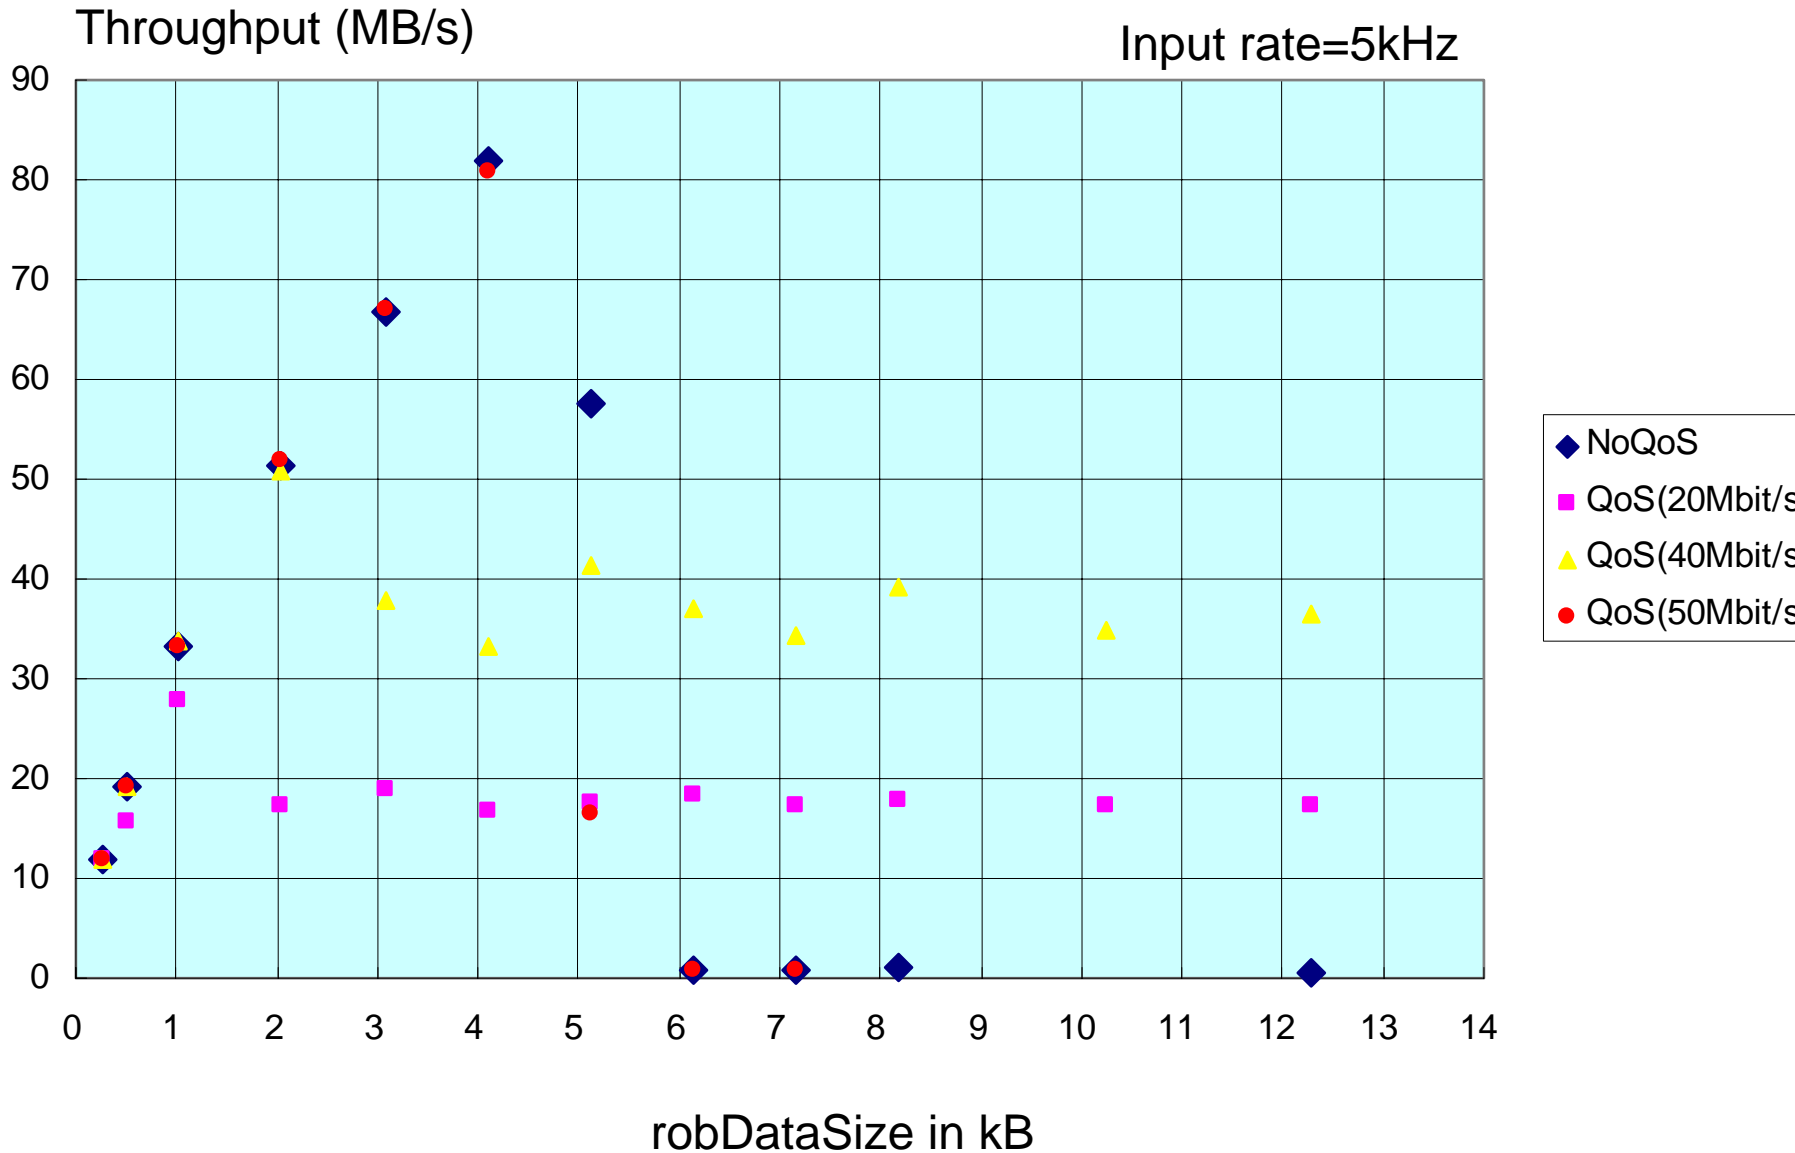
Throughput without QoS



Results from the 6x1 system without QoS

1. Performed push and pull scenarios with similar setups.
2. Plateau observed in EoE rates after it reaches peak rate of 5kHz in push scenario and 4kHz in pull scenario when the robDataSize was 1kB.
3. Throughput reached over 90MB/s in pull scenario and 80MB/s in push scenario, but the throughput in push scenario dropped when robDataSize increased.
 - Lots of re-asks observed

Throughput with QoS in push scenario



Results from the 6x1 system with QoS

In push scenario

1. Plateau observed in the throughput after it reaches the QoS assigned rate at 40Mbit/s and 20Mbit/s. In the plateau, re-asks decreased while there were lots of re-asks without QoS.
2. The assigned rate of 50Mbit/s did not work.

Summary

1. The results from performance of DC software at CERN were consistent to that from ICEPP. See the talk of result from ICEPP at Atlas week, Feb 2003.
2. QoS improved the event builder performance in push scenario.

Plan

1. QoS performance measurement on 24 PCs.
2. QoS performance measurement in pull scenario on N x M system.

Acknowledgments

Thanks to Stefan, Marian and other staff at Bldg.513
for their support

Thanks to Hanspeter, Christian and other DC staff for
their help

Thanks to Andreas and Per for sharing Level2
testbed at Bldg.513