FUJITSU

shaping tomorrow with you

Understanding Storage Traffic Characteristics on Enterprise Virtual Desktop Infrastructure

Chunghan Lee, Tatsuo Kumano, Tatsuma Matsuki,
Hiroshi Endo, Naoto Fukumoto, Mariko Sugawara

ACM SYSTOR 2017

05.23.2017

Fujitsu Laboratories Ltd.

Introduction

- Virtual desktop infrastructure (VDI)
  - It can enhance security and reduce management costs
  - It has been used in real-world offices and universities

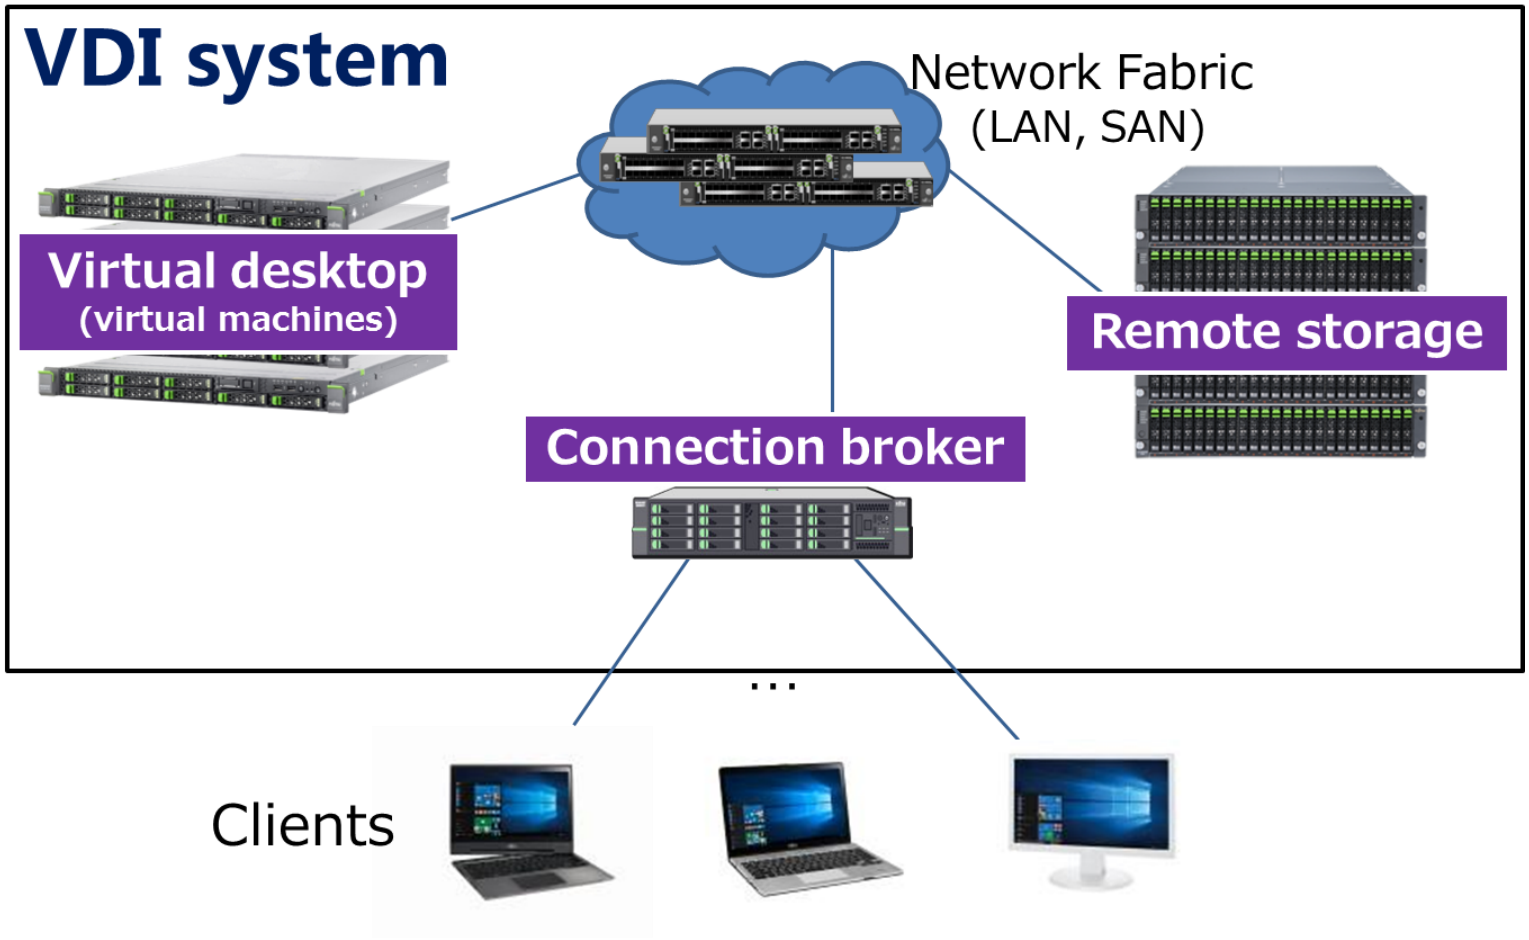

Previous works on VDI systems

FUJITSU

- Previous studies have analyzed the common Internet file system (CIFS) and storage traffic in personal cloud storage

- Few prior works on VDI system
  - Only one field study [FAST'11] with a small-scale VDI (55 VMs)
  - They analyzed storage traffic characteristics on the small-scale VDI by collecting traces from a specific component (hypervisor)

Research goal

FUJITSU

A deep understanding of storage traffic characteristics: FC traffic and VM behavior with large-scale VDI

Approach

FUJITSU

- **Focus on two perspectives: fibre channel (FC) traffic and VM behavior**
  - Gather actual data from a part of an enterprise VDI for 28 consecutive days
  - Datasets consist of various types of traces, including a usage questionnaire and active and passive measurements
  - Analyze datasets corresponding to 79.8 TiB of storage traffic on approximately 300 VMs

Key findings

FUJITSU

- **We mainly present key findings (3,4 and 6)**

FC traffic	1. Read traffic is dominant
	2. During an update storm, a large number of transactions access similar address blocks in bursts
	3. During the storm, the write response times of large size transactions are degraded
VM behavior	4. Although few VMs generate dominant traffic, their impact is small
	5. Anti-virus and Windows service in idle VMs generate large traffic
	6. All active VMs generate only 25% of traffic

Refer to our detailed contributions in the paper ☺

Target VDI

FUJITSU

- Overview of target VDI
  - It is a part of a office VDI system
    - Six servers (300 VMs) and six block storage with FCoE fabric
  - It provides a designated Windows VM to each user
  - It also provides pre-installed applications, such as a mailer (Outlook), Microsoft Office, and anti-virus software (McAfee)
  - Users can install desired applications and save their data anytime

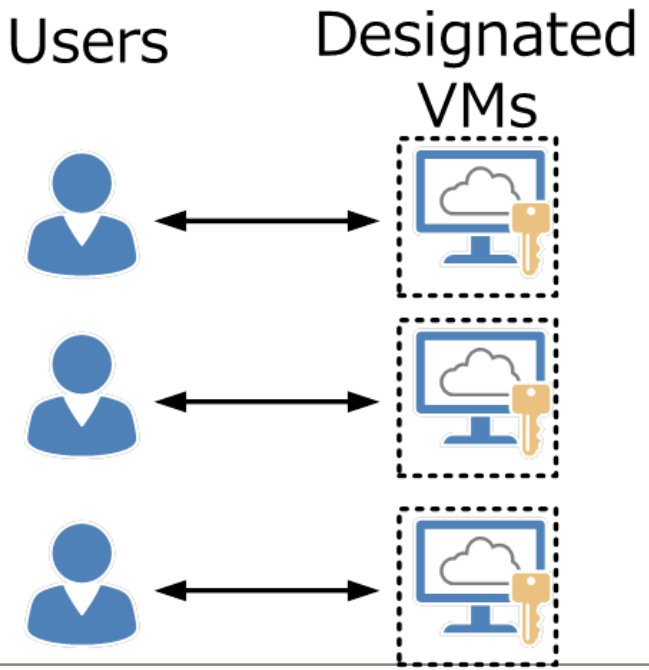

How to measure target VDI?

FUJITSU

1. We cannot install monitoring components on a VM or hypervisor

→ We captured FC traffic, and gathered performance metrics from the storage and VDI servers

2. The VMs are always running to reduce the impact of time-consuming operations, such as VM restart

→ We gathered connection broker logs from VDI portal to distinguish active VMs (connected to users)

3. We cannot log on to user VMs, and it is hard to inspect user applications on the VM

→ We gathered a list of applications at all VMs, and distributed a usage questionnaire to determine which applications were actually run

Measurement methodology

- Gather various datasets for 28 consecutive days

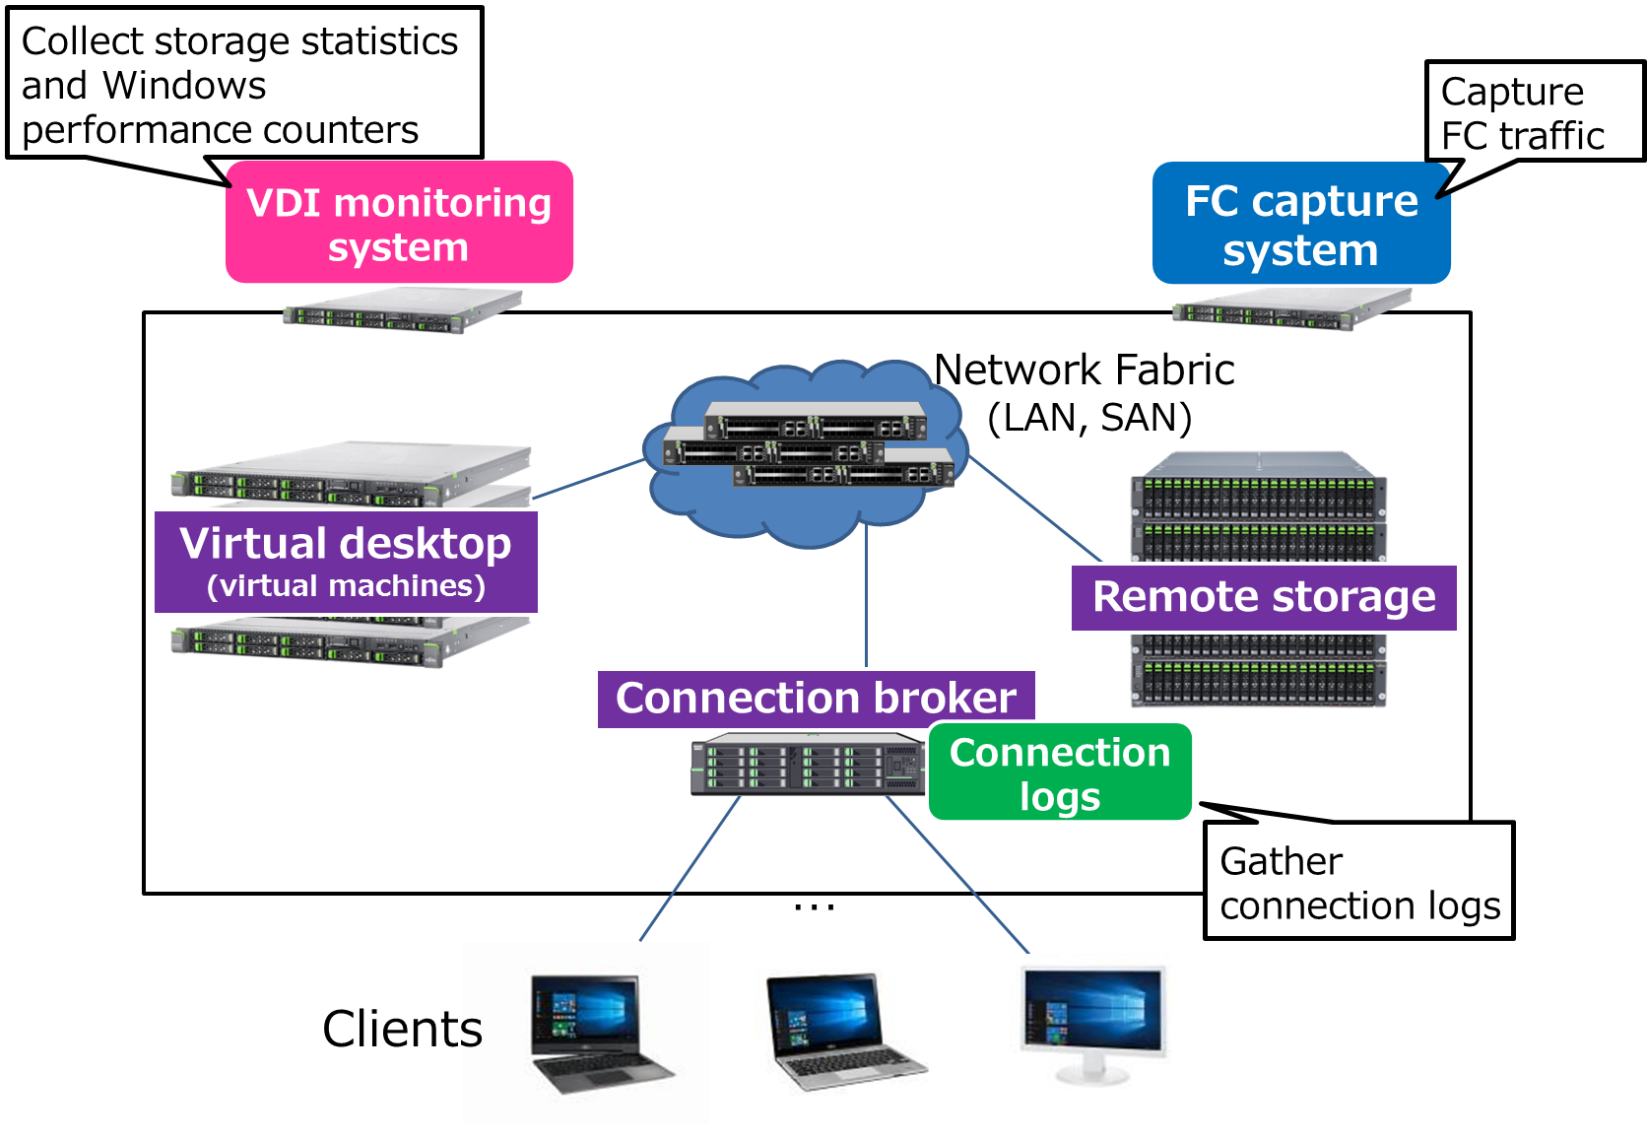

Datasets for analysis

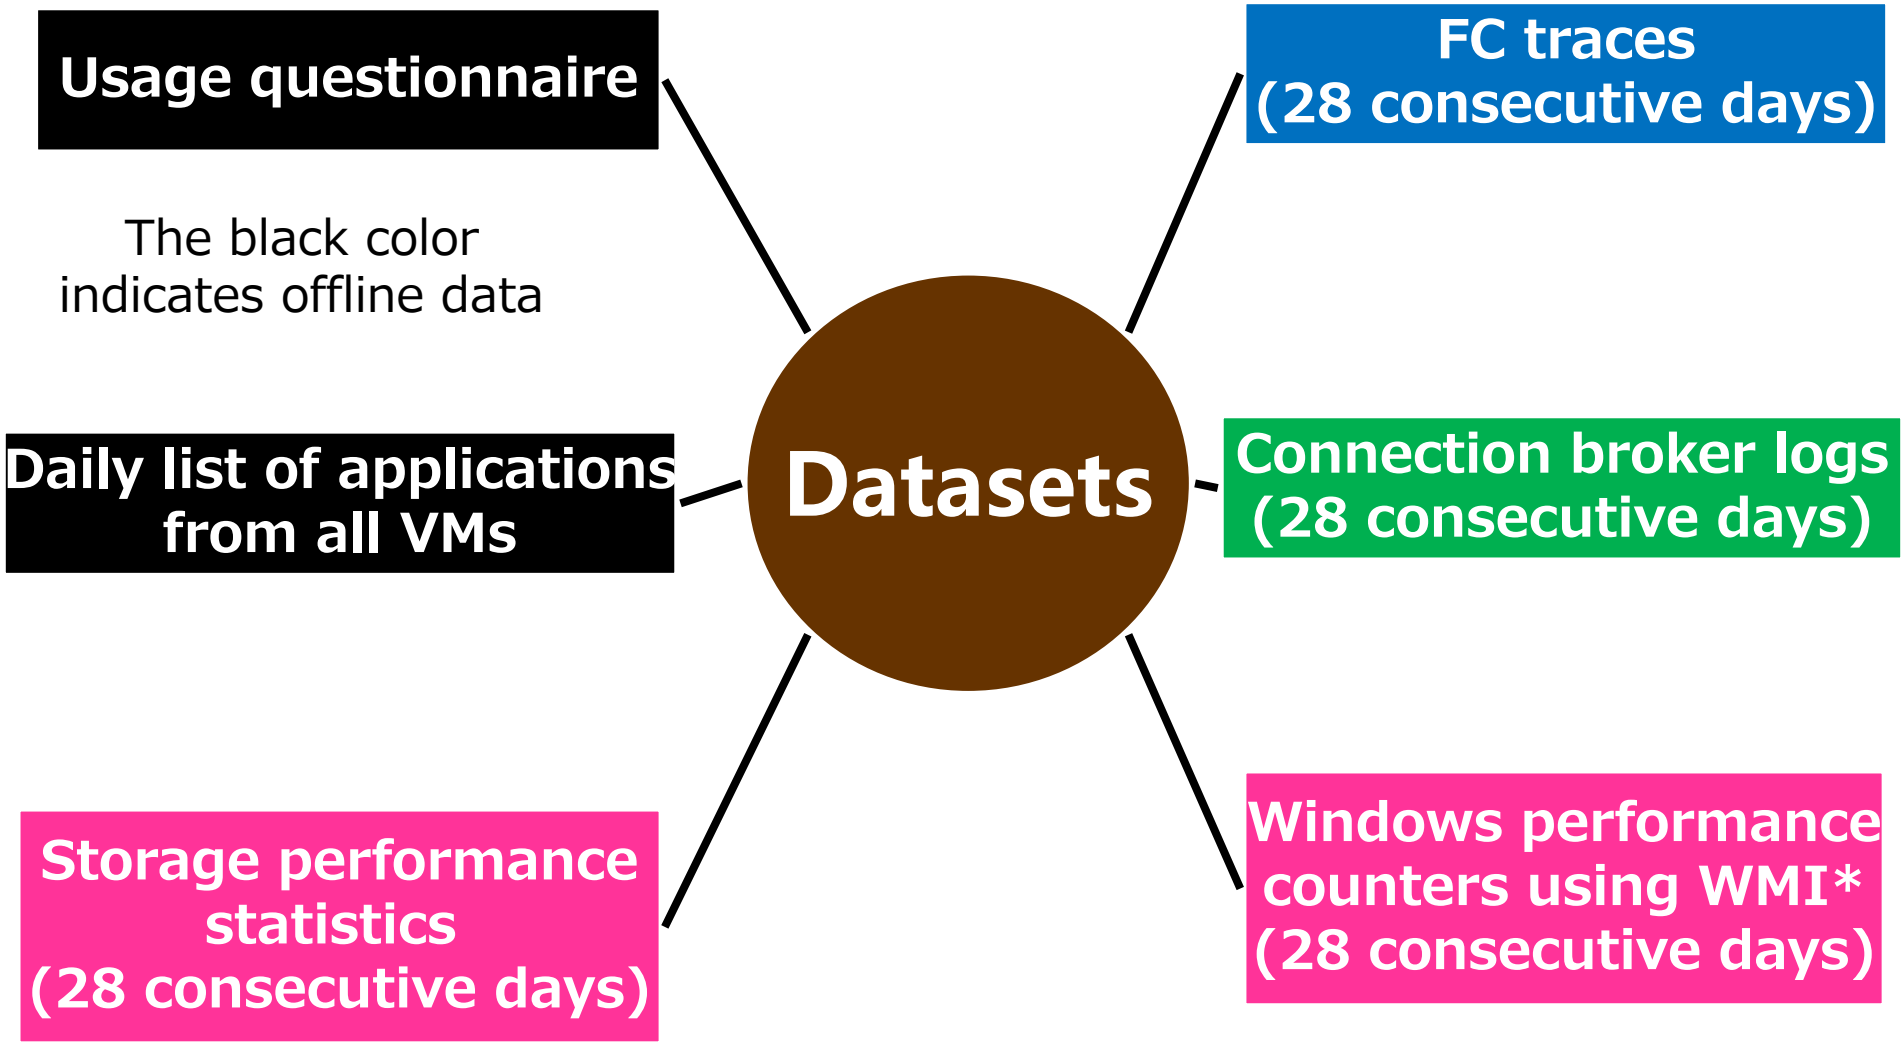

*Windows Management Instrumentation (WMI)

Applications installed on the VMs

FUJITSU

- TreeMap
  - We manually and empirically classified the applications into nine categories
  - The pre-installed software, such as Microsoft Office and anti-virus, had the large proportion
  - We also found the similar characteristics with the usage questionnaire results
  - The applications were similar to those of the small-scale VDI [FAST'11]

Deeper colors and larger rectangles indicate larger proportions

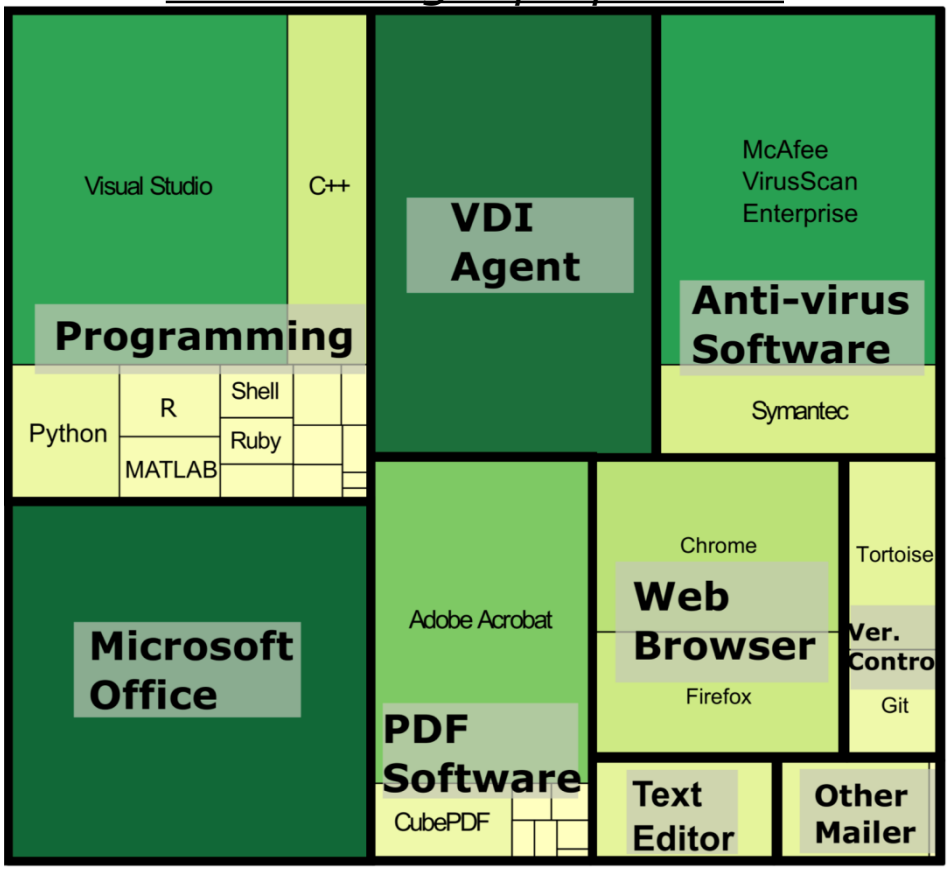

VM-level traffic

- Read and write traffic
  - Burst read traffic (deep black color) occurs, while continuous small write traffic (light gray color) is observed
    - Read traffic is dominant (82%)
  - Massive traffic is intensively generated in the morning (09:00)
  - The VM-level traffic widely fluctuates and depends both on time and the VMs
    - It would be possible to exist *heavy VMs* that generated dominant traffic

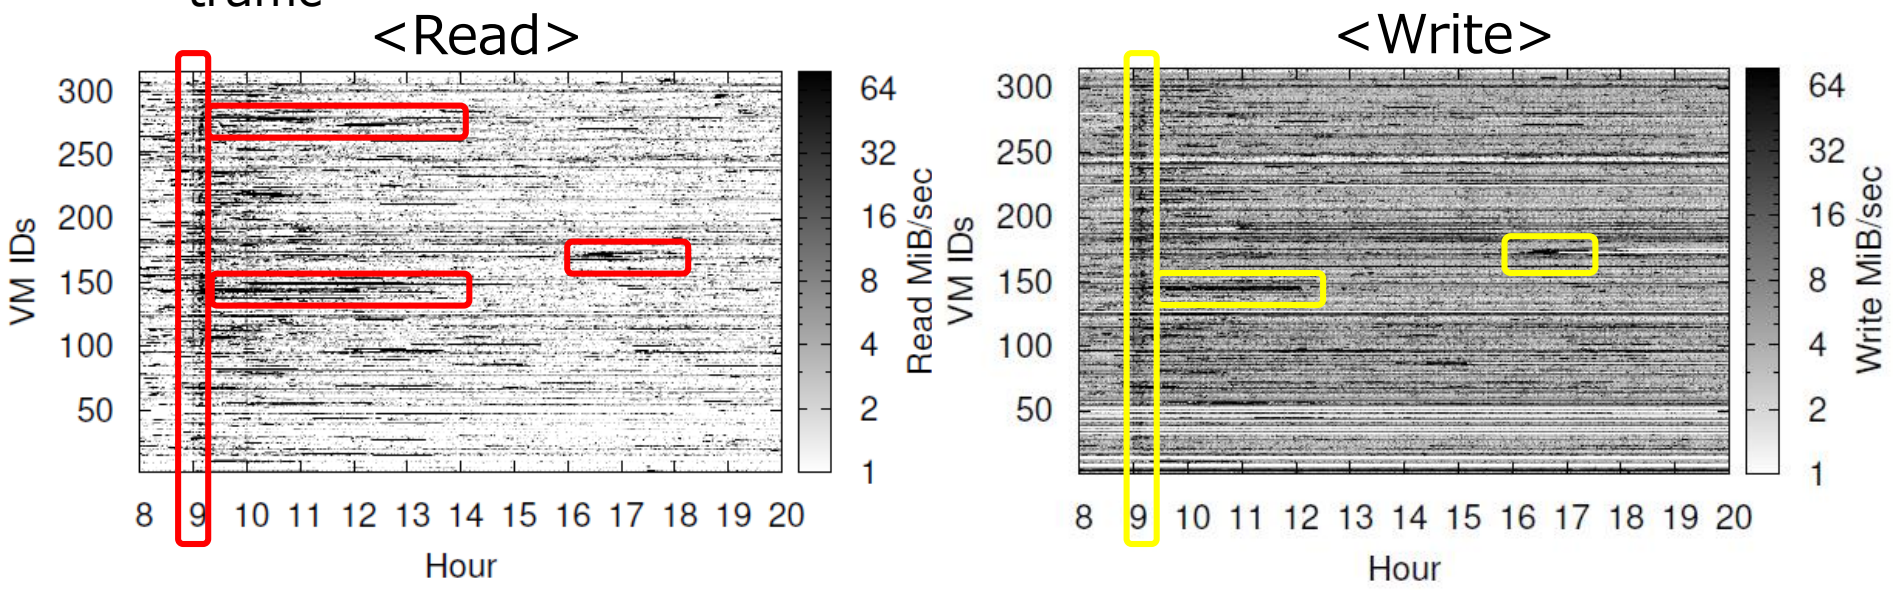

Questions from the results

FUJITSU

1. What happened in the morning (09:00-10:00)?

2. Why read traffic is dominant?

3. How the heavy VMs are affected to the VDI system?

What happened in the morning?

- Daily distribution of write throughput
  - From application logs, we found the regular update of anti-virus
  - The update was simultaneously executed by hundreds of VMs, and it acts as an **update storm**
  - During the update storm, the large throughput of both read and write traffic was achieved, and their response time was also largely increased

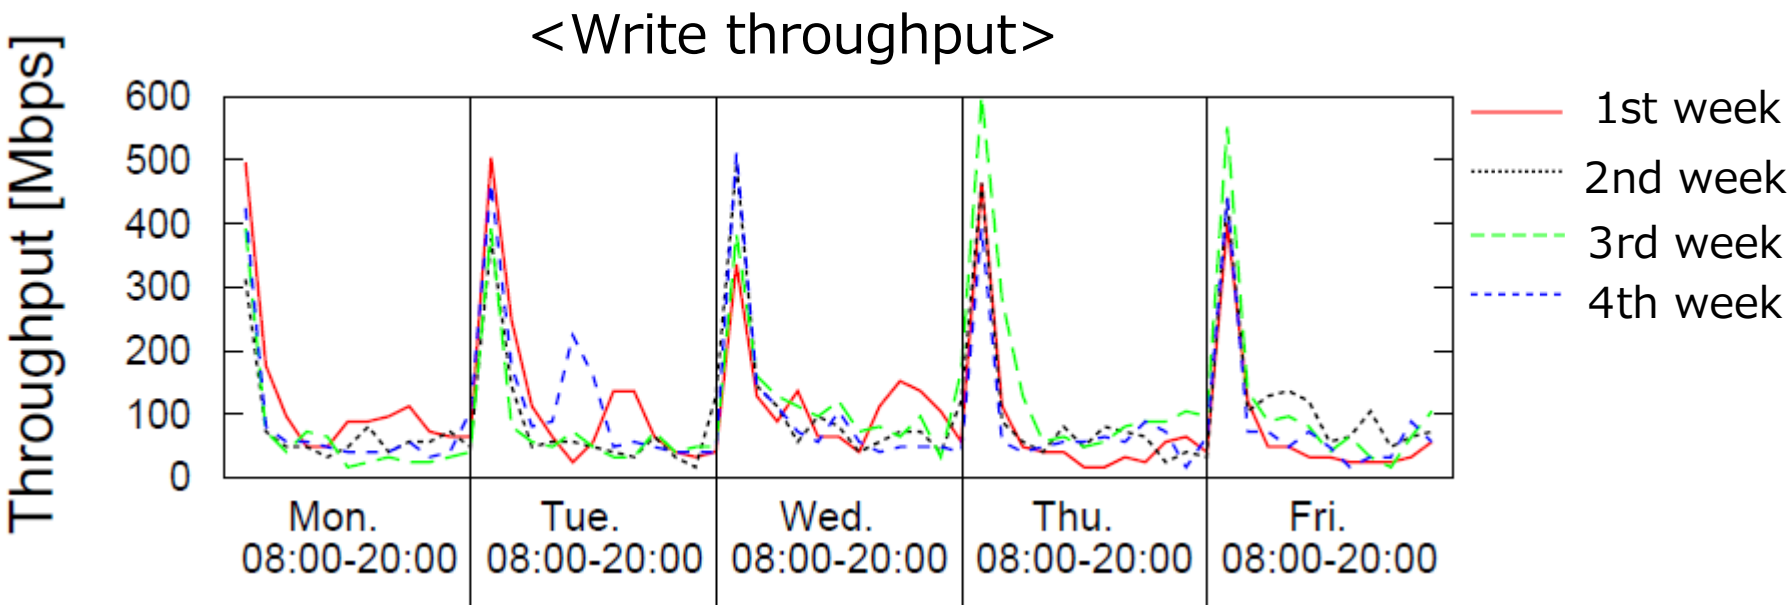

Characteristics of update storm – (1)

FUJITSU

- The storage response time was largely degraded from 20 ms during the update storm
  - 20 ms : the recommend response time by VMWare
- Write response time
  - In the morning, the CDF shape differs from the other time bins
- Write size distribution
  - In the morning, the large size (128 KiB) has large proportion

<Response time over 20ms>

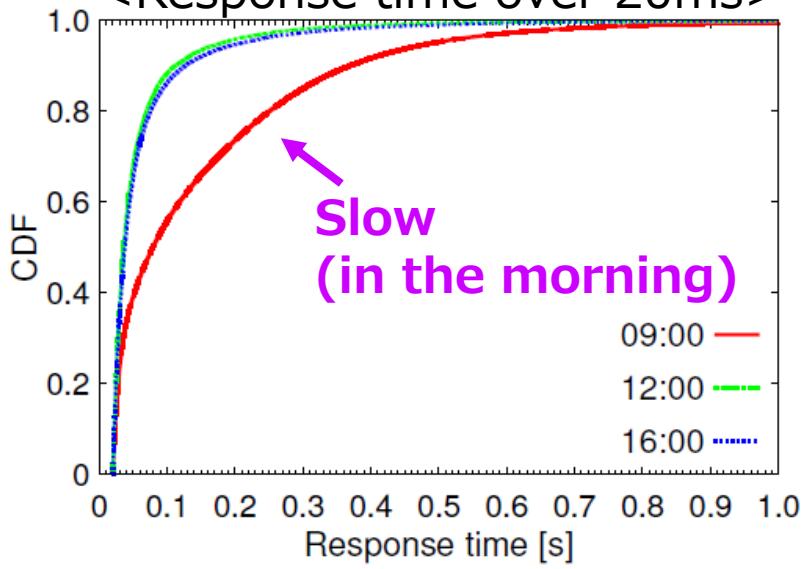

<Size distribution over 20ms>

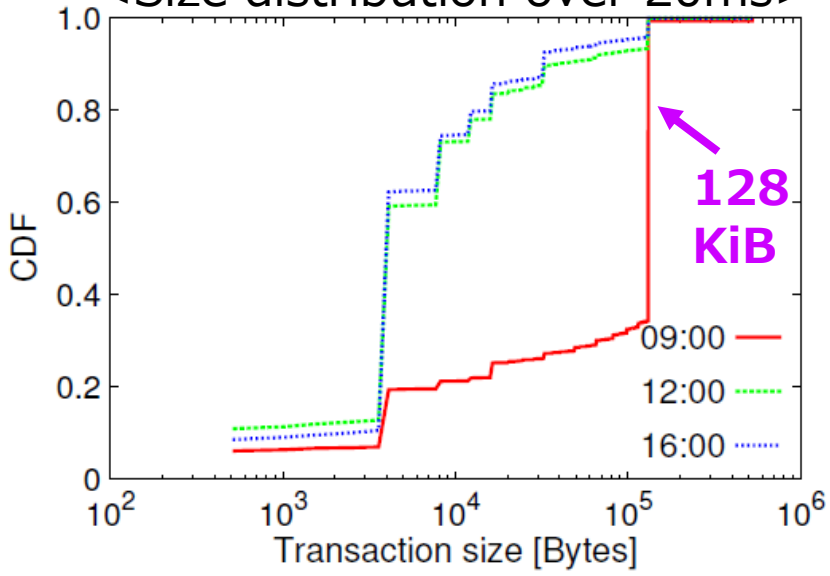

** VMware. Storage I/O control technical overview and considerations for deployment http://www.vmware.com/files/pdf/techpaper/VMW-vSphere41-SIOC.pdf.

Characteristics of update storm – (2)

FUJITSU

- Write response time is affected by write transaction size
  - The write response time of large size transaction is greatly affected even when write cache hit rate is slightly decreased
  - Under a lack of write cache memory, the large write transactions are much more affected than the small ones

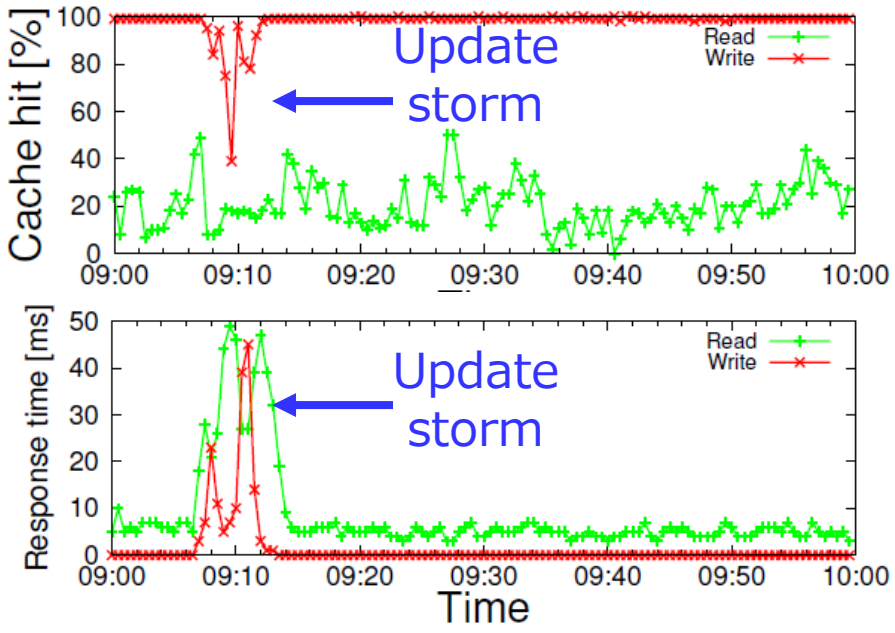

<Write size distribution over 20ms>

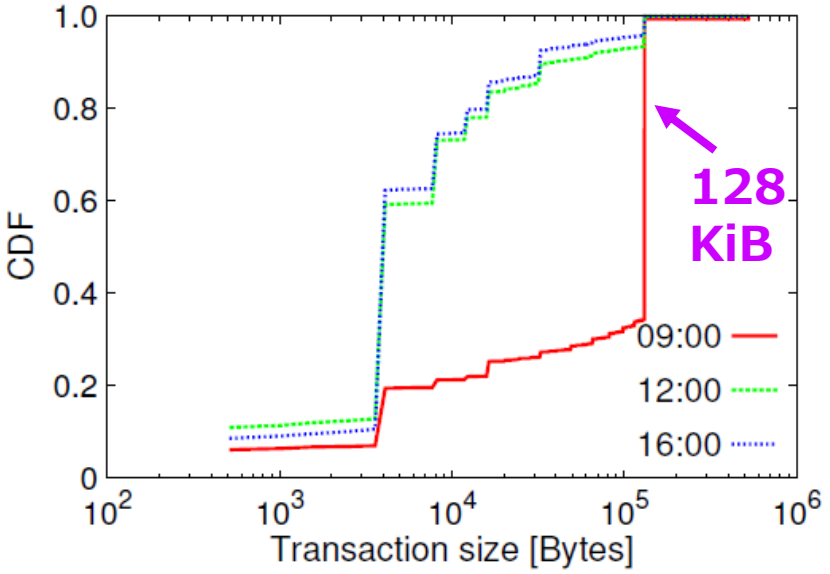

Characteristics of update storm – (3)

FUJITSU

- Implication from the analysis results

The transaction size for write cache hit rate should be considered to improve overall VDI performance

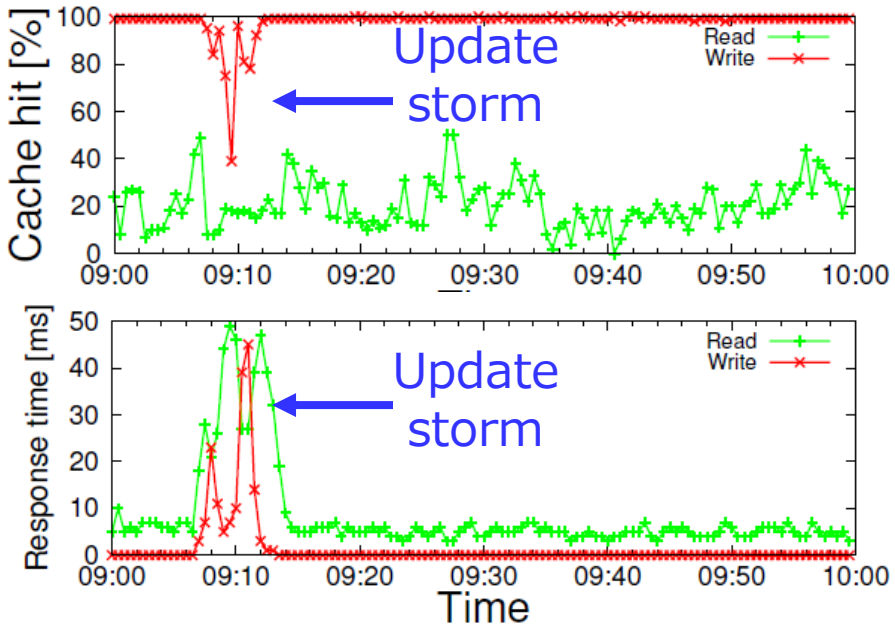

<Write size distribution over 20ms>

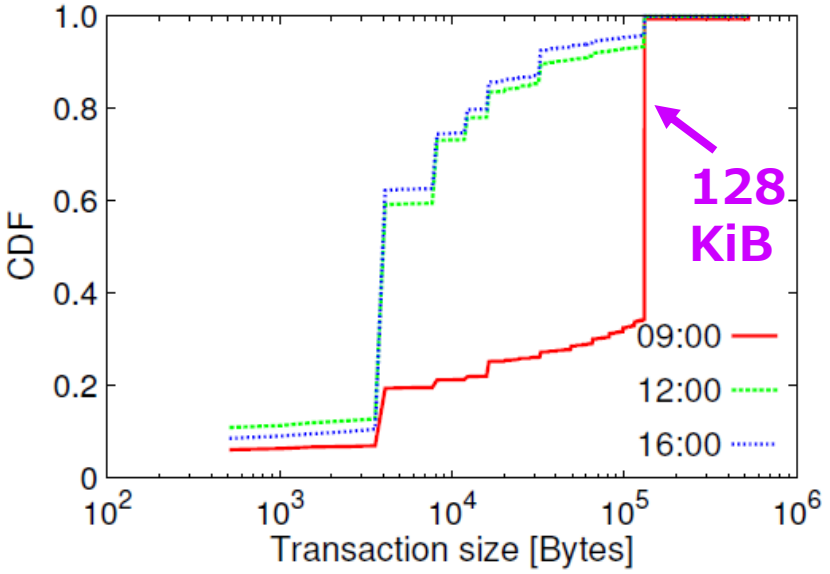

Major cause of dominant read traffic – (1)

- The idle VMs generate large read traffic
  - The idle VMs (not connected to users) account for 75.5% of total traffic
    - Anti-virus and Windows service were major applications to generate large read traffic
  - The active VMs (connected to users) issue 24.5% of total traffic

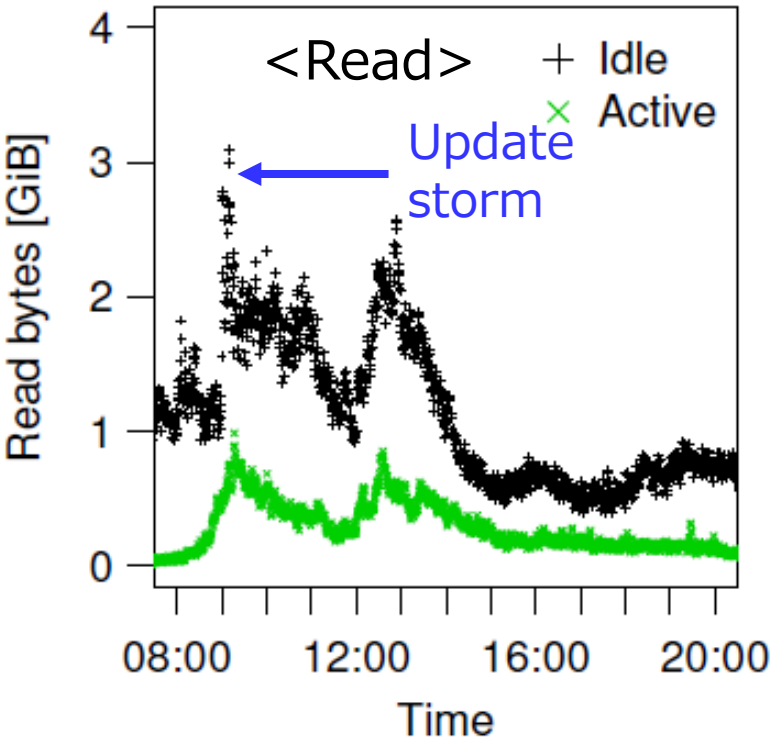

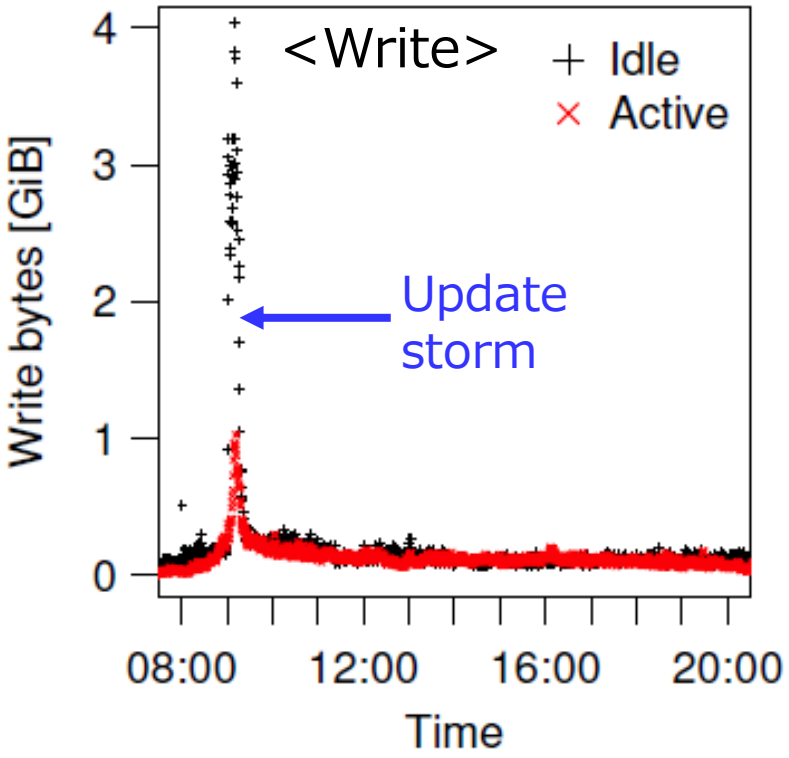

Major cause of dominant read traffic – (2)

- Implication from the analysis results

Traffic from active VMs should be more highly prioritized to shorten their response time

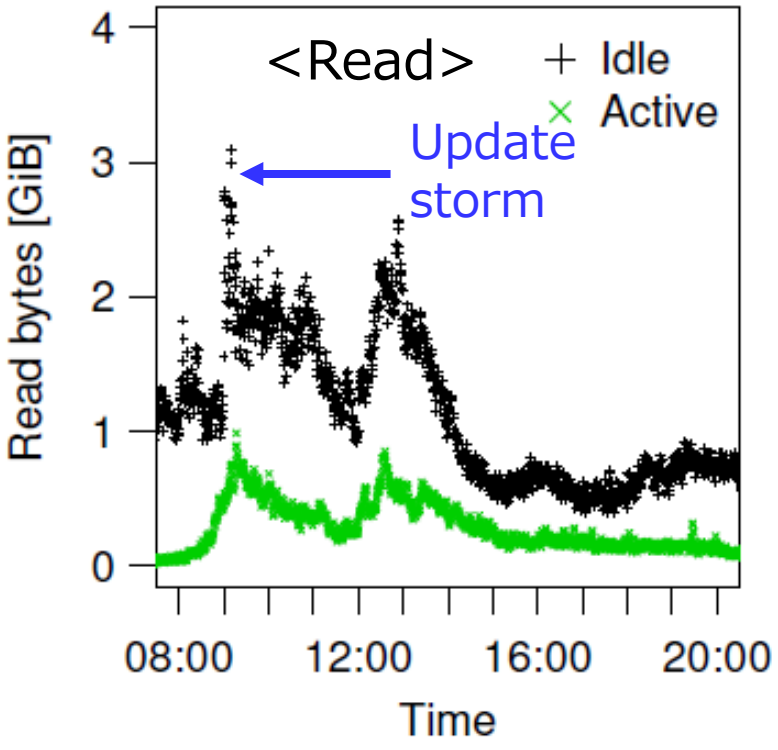

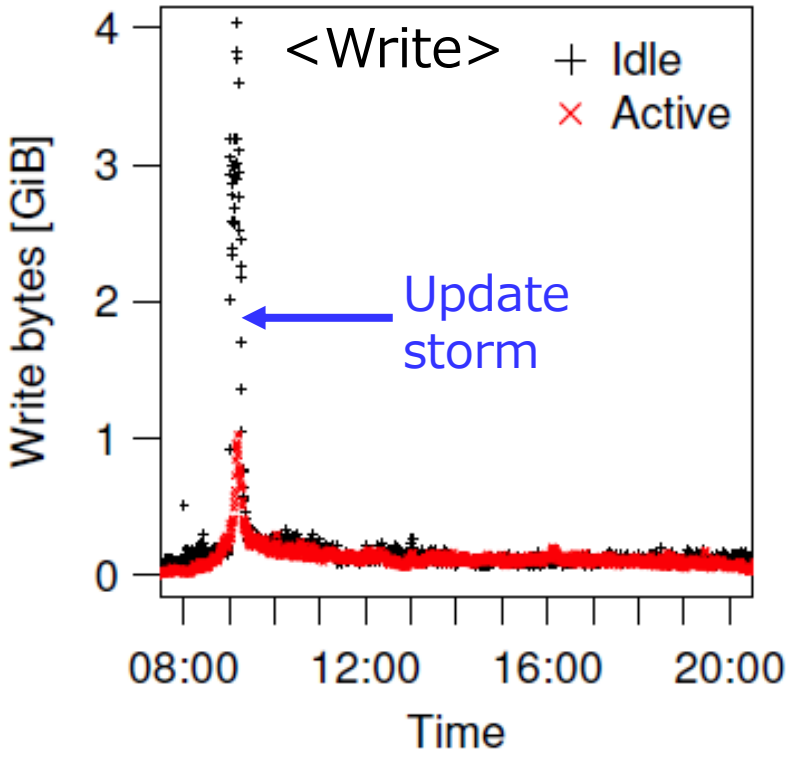

Traffic distribution across VMs – (1)

- The morning (09:00-10:00)
  - The traffic distribution at all VMs is almost homogeneous traffic
  - The update storm is a major cause
- The lunch time (12:00-13:00)
  - The traffic distribution at all VMs is heterogeneous traffic
  - The scheduled virus scan by some VMs is a major cause
  - 1% of active VMs only generate 16.6% of traffic

<Gini coefficient>

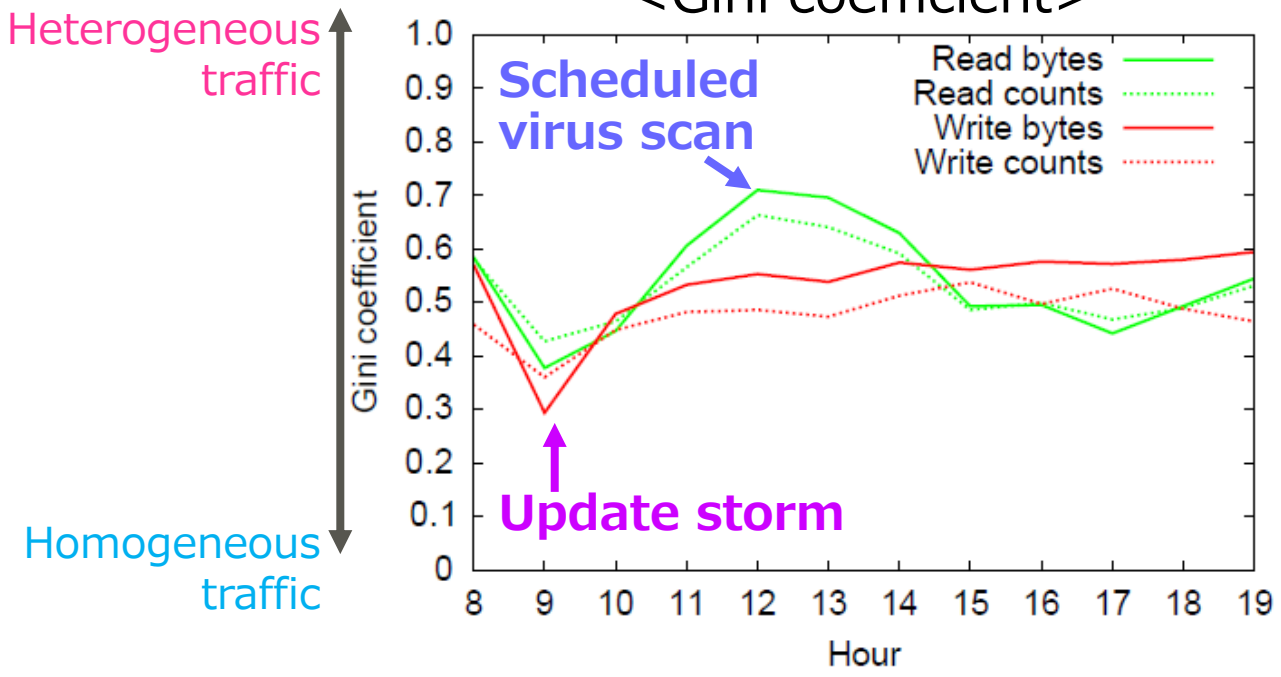

<Lorenz curve of read bytes>

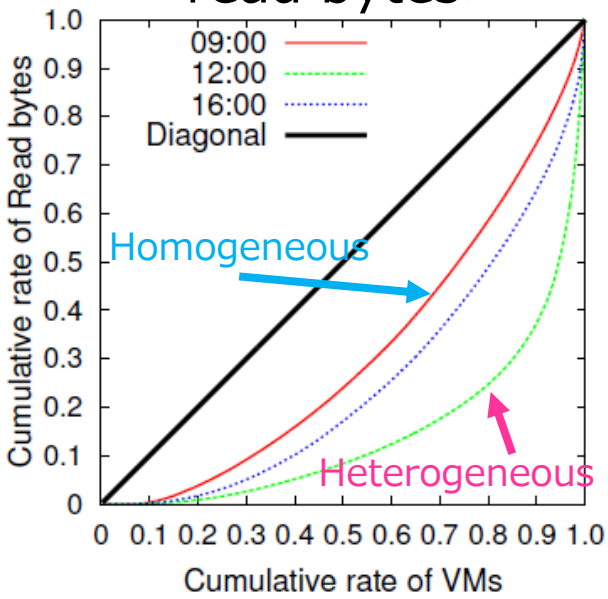

Traffic distribution across VMs – (2)

FUJITSU

- Implication from the analysis results

The impact of heavy VMs is small and especial load balancing strategy is not needed to office VDI system

Heterogeneous traffic

Homogeneous traffic

<Gini coefficient>

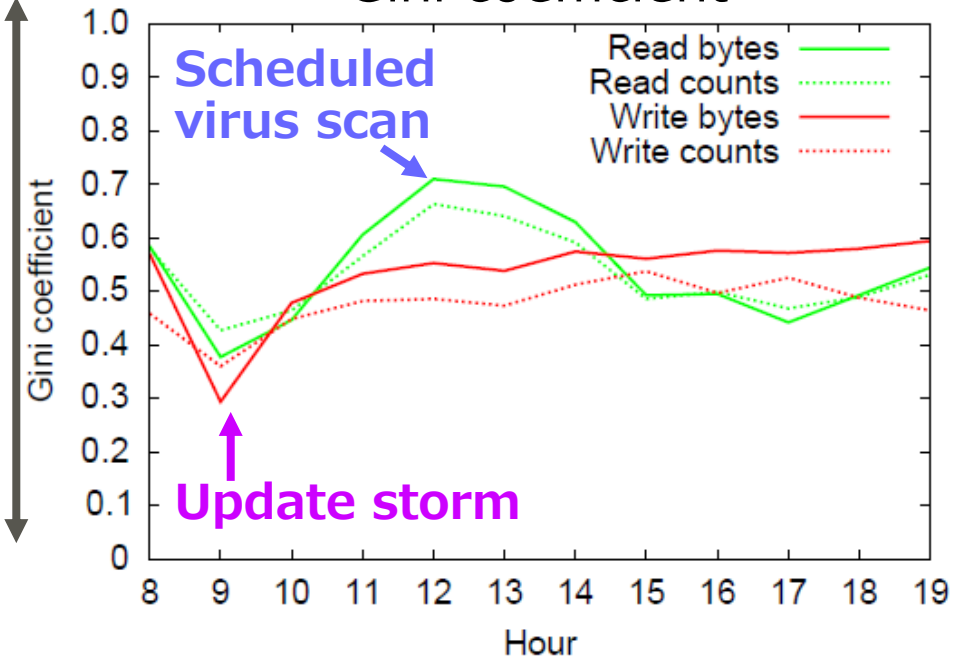

<Lorenz curve of read bytes>

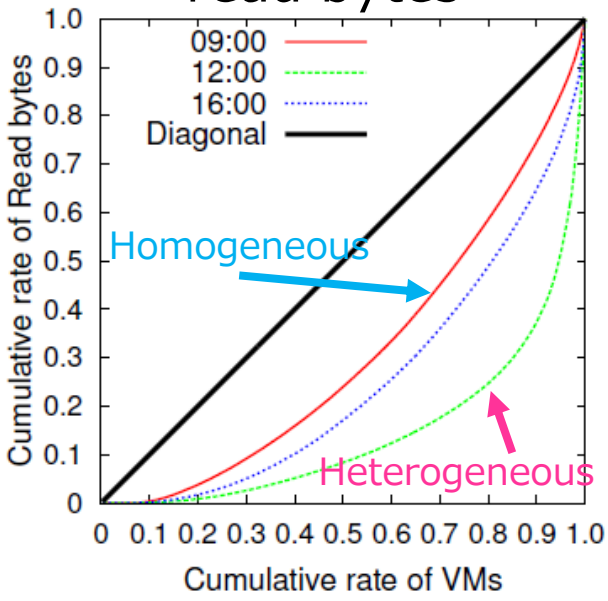

Conclusion

FUJITSU

1. What happened in the morning (09:00-10:00)?

→ **The update storm of anti-virus software**

The transaction size for write cache hit rate should be considered to improve overall VDI performance

2. Why read traffic is dominant?

→ **The heavy read traffic from idle VMs**

Traffic from active VMs should be more highly prioritized to shorten their response time

3. How the heavy VMs are affected to the VDI system?

→ **The impact of heavy VMs is small and 1% of active VMs generate only 16.6% of traffic**

The VDI system can perform sufficiently by using uniform load balancing strategy

Future work

FUJITSU

- Analyze the characteristics of user application workloads on the active VMs
- Investigate the storage traffic characteristics on different types of VDI system

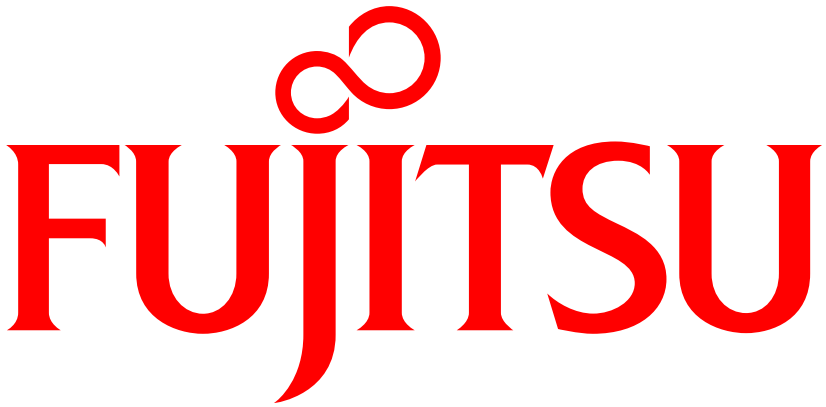

shaping tomorrow with you

Daily active users

FUJITSU

- Patterns of daily active users
  - Two patterns are observed: weekday and weekend
  - We determine the measurement period (08:00-20:00) based on daily active users

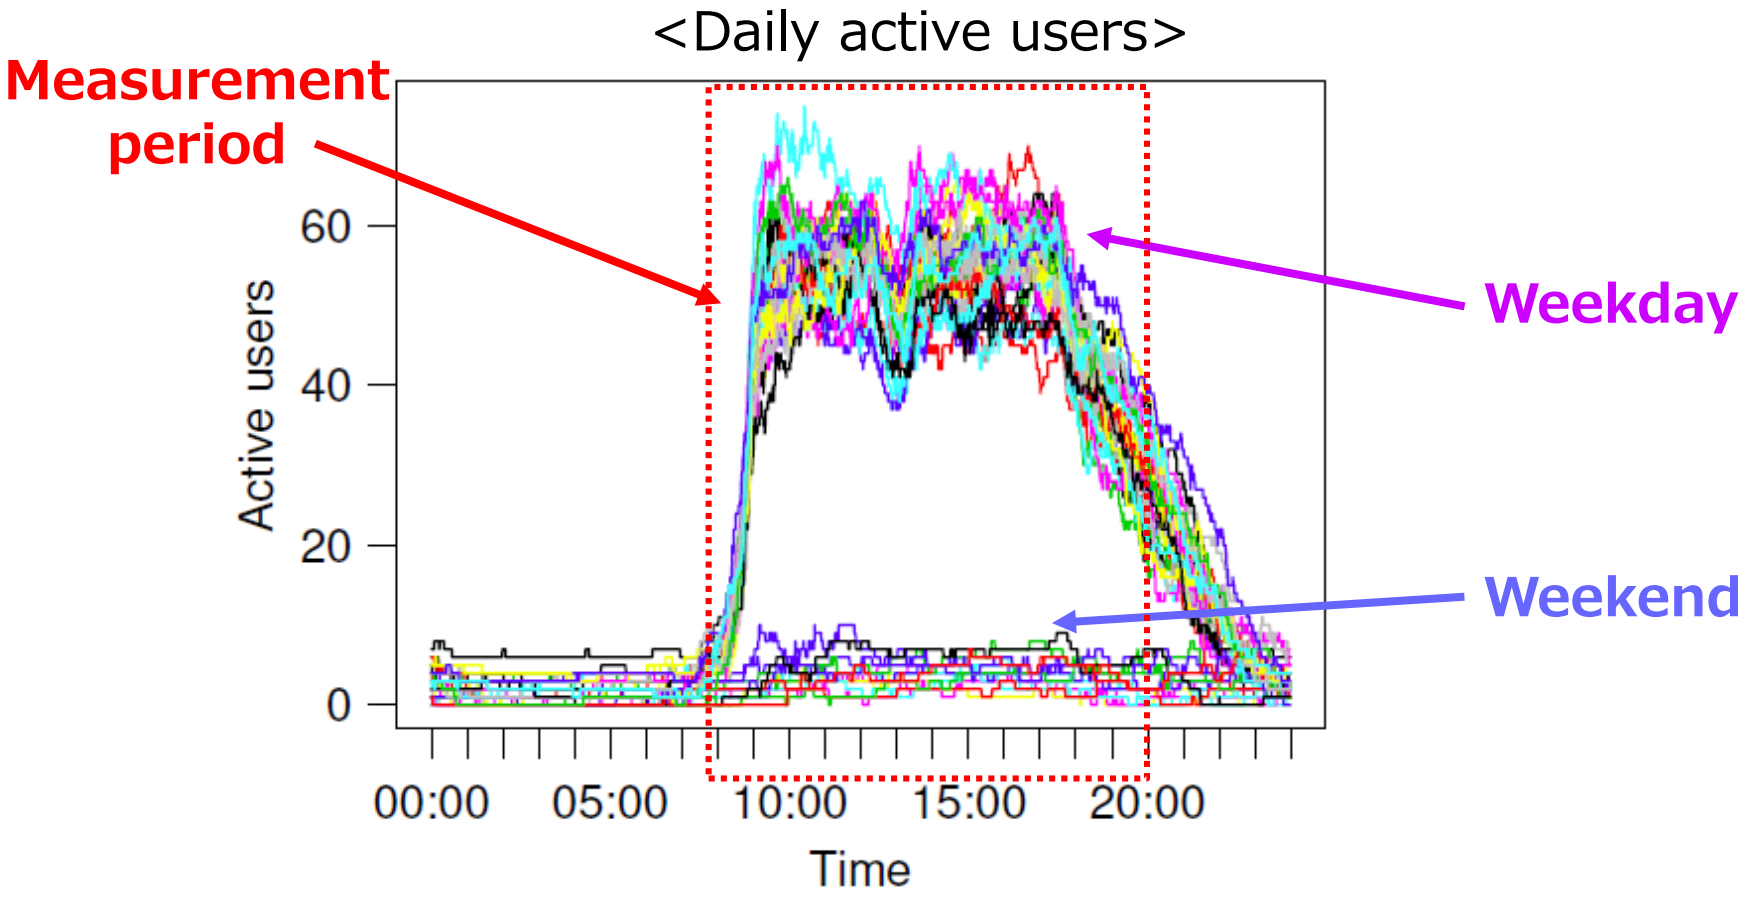

Major applications on idle VMs

FUJITSU

- Windows service, anti-virus, and system-profiling are major applications to generate large read traffic
- The read traffic at idle VM is 93%

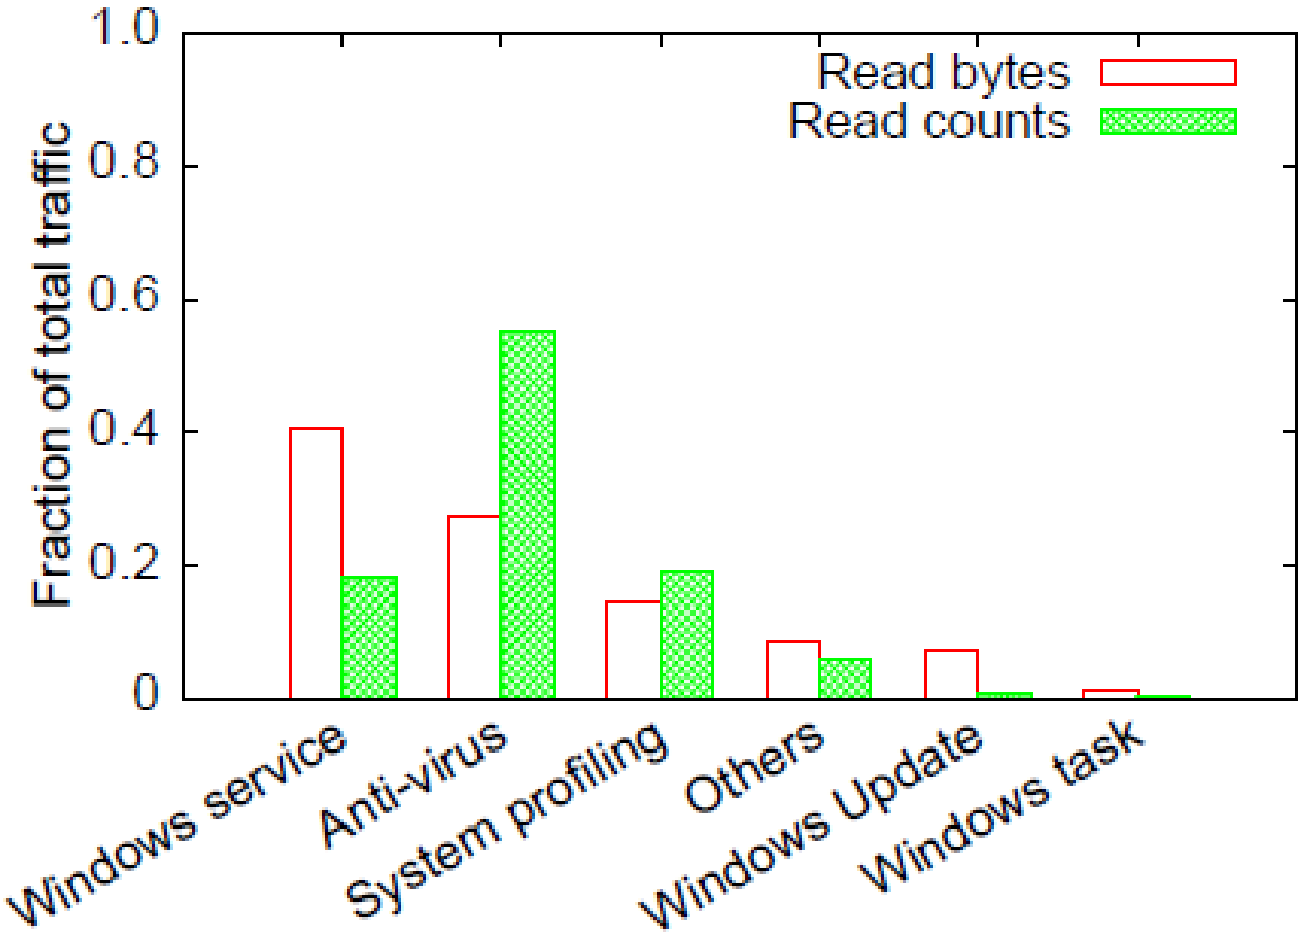

Daily read throughput

- Obvious periodicity pattern with three time bins
  - Morning : 09:00-10:00
  - Lunch time : 12:00-13:00
  - Afternoon : 16:00-17:00

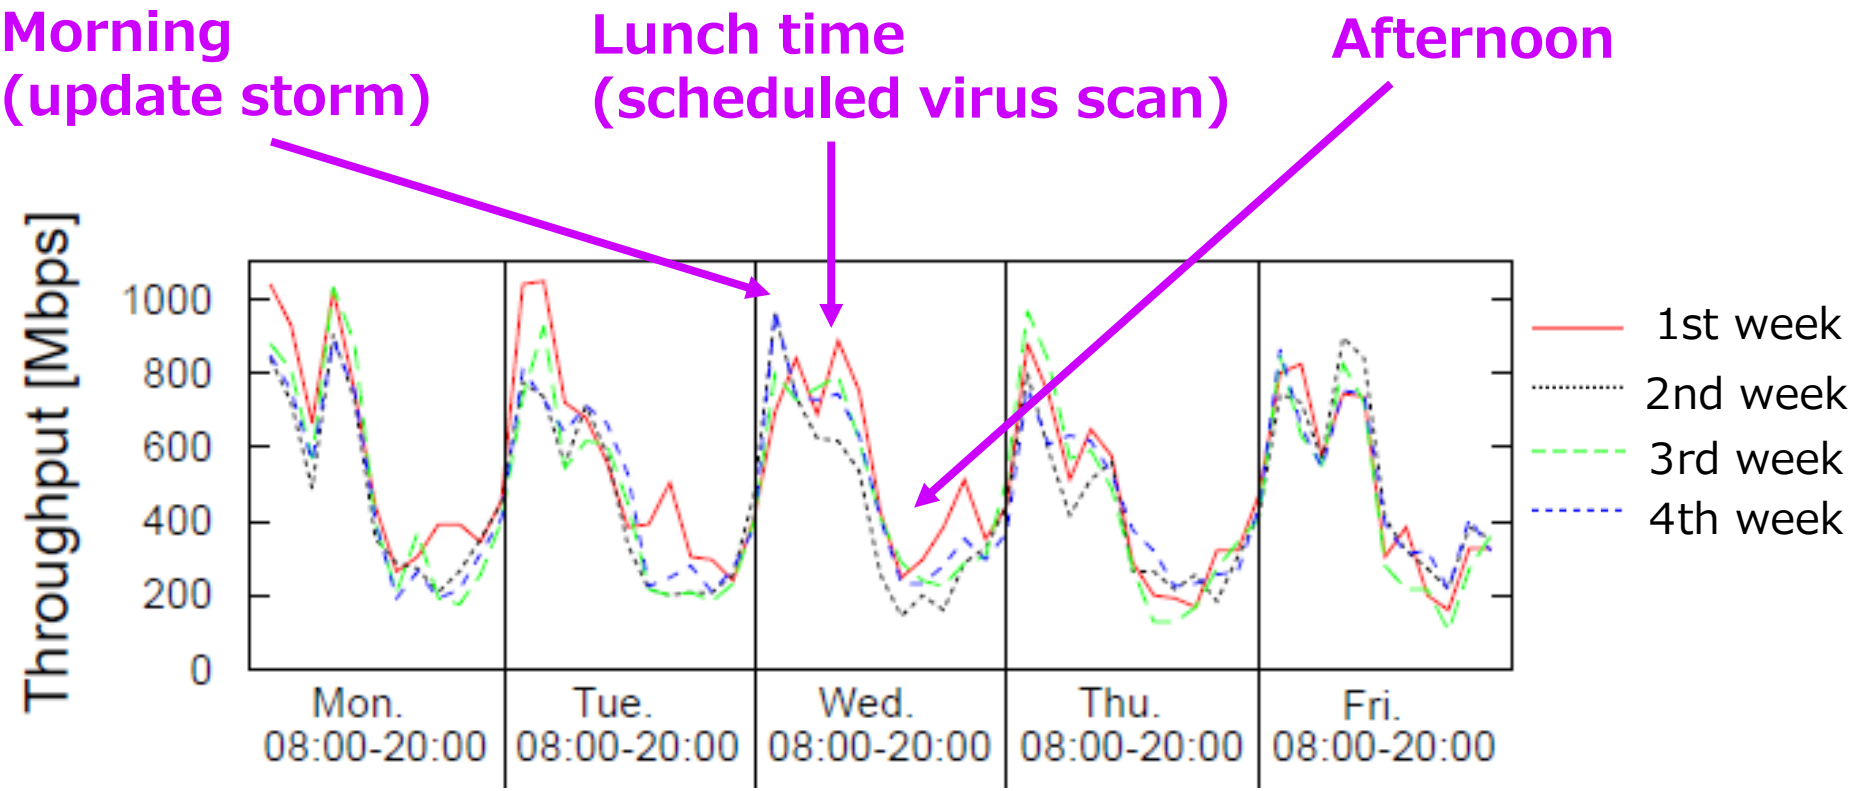

Spatial distribution at LUN

- During the update storm, a large number of transactions access similar address blocks in bursts over a short period of time

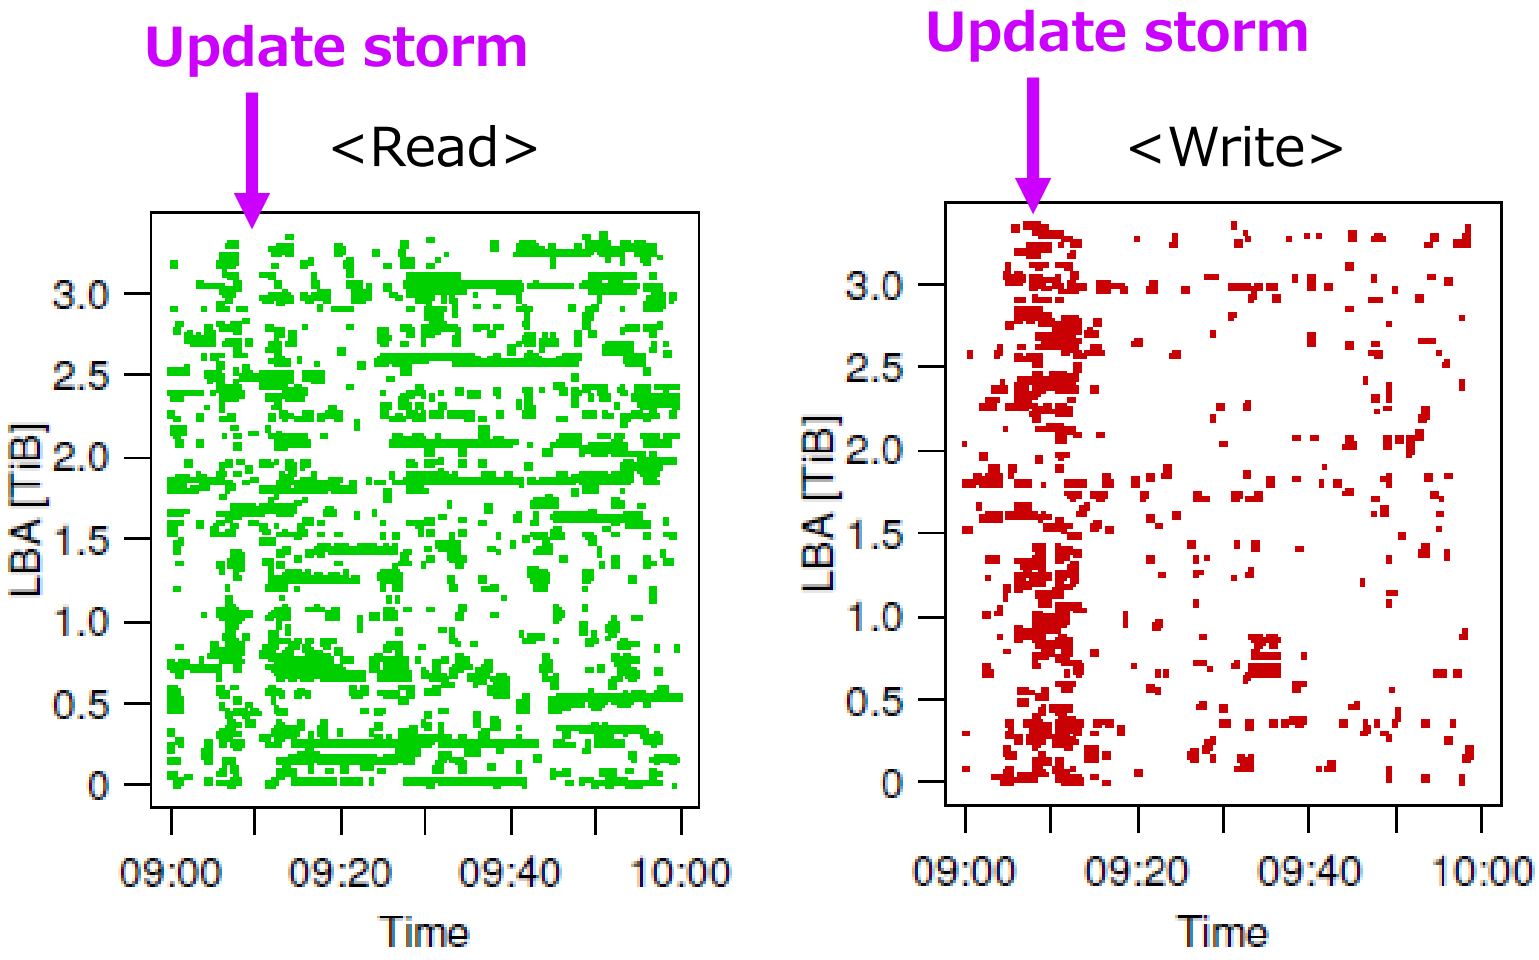

Target VDI

FUJITSU

- Using designated VM
  - A designated VM is provided to each user
  - Users log on their VM only, install desired applications and save their data anytime

Our target VDI

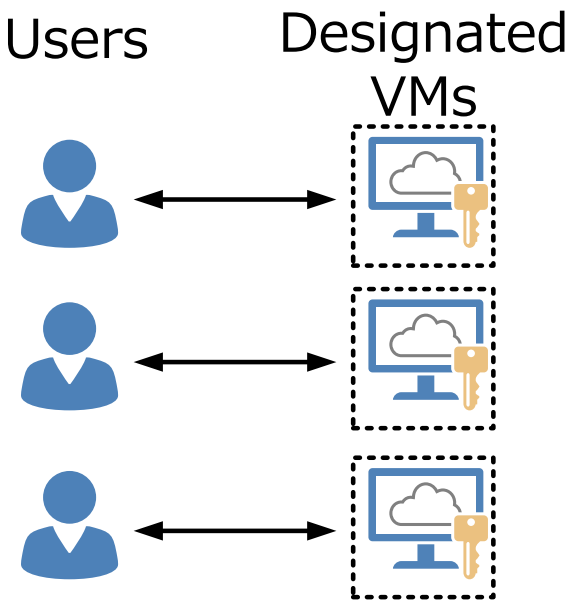

- Pool-based VDI
  - A new VM is always provided to each user
  - Users cannot install desired applications because they have no permission to control their VMs
  - If a connection is broken to a VM, all of status are deleted