

IT VIRTUALIZATION FOR DISASTER MITIGATION AND RECOVERY

Maurício Tsugawa
Renato Figueiredo
José Fortes



Takahiro Hirofuchi
Hidemoto Nakada
Ryousei Takano



J-RAPID Symposium
Sendai, March 6-7th, 2013



Motivation

- Information Technology (IT) is applied in almost all infrastructures and services
 - IT services needs to be quickly recovered from damages
 - Desirably keep IT services undisturbed
- Most catastrophic events cannot be predicted
 - Typical disaster recovery (DR) services are expensive
 - Applications need to be adapted for DR
- Opportunities
 - Emerging trend to use cloud services for disaster recovery
 - Use virtualization technologies to enable resilient IT services
 - VMs are movable
 - On-demand migration
 - Lower cost

Project Overview

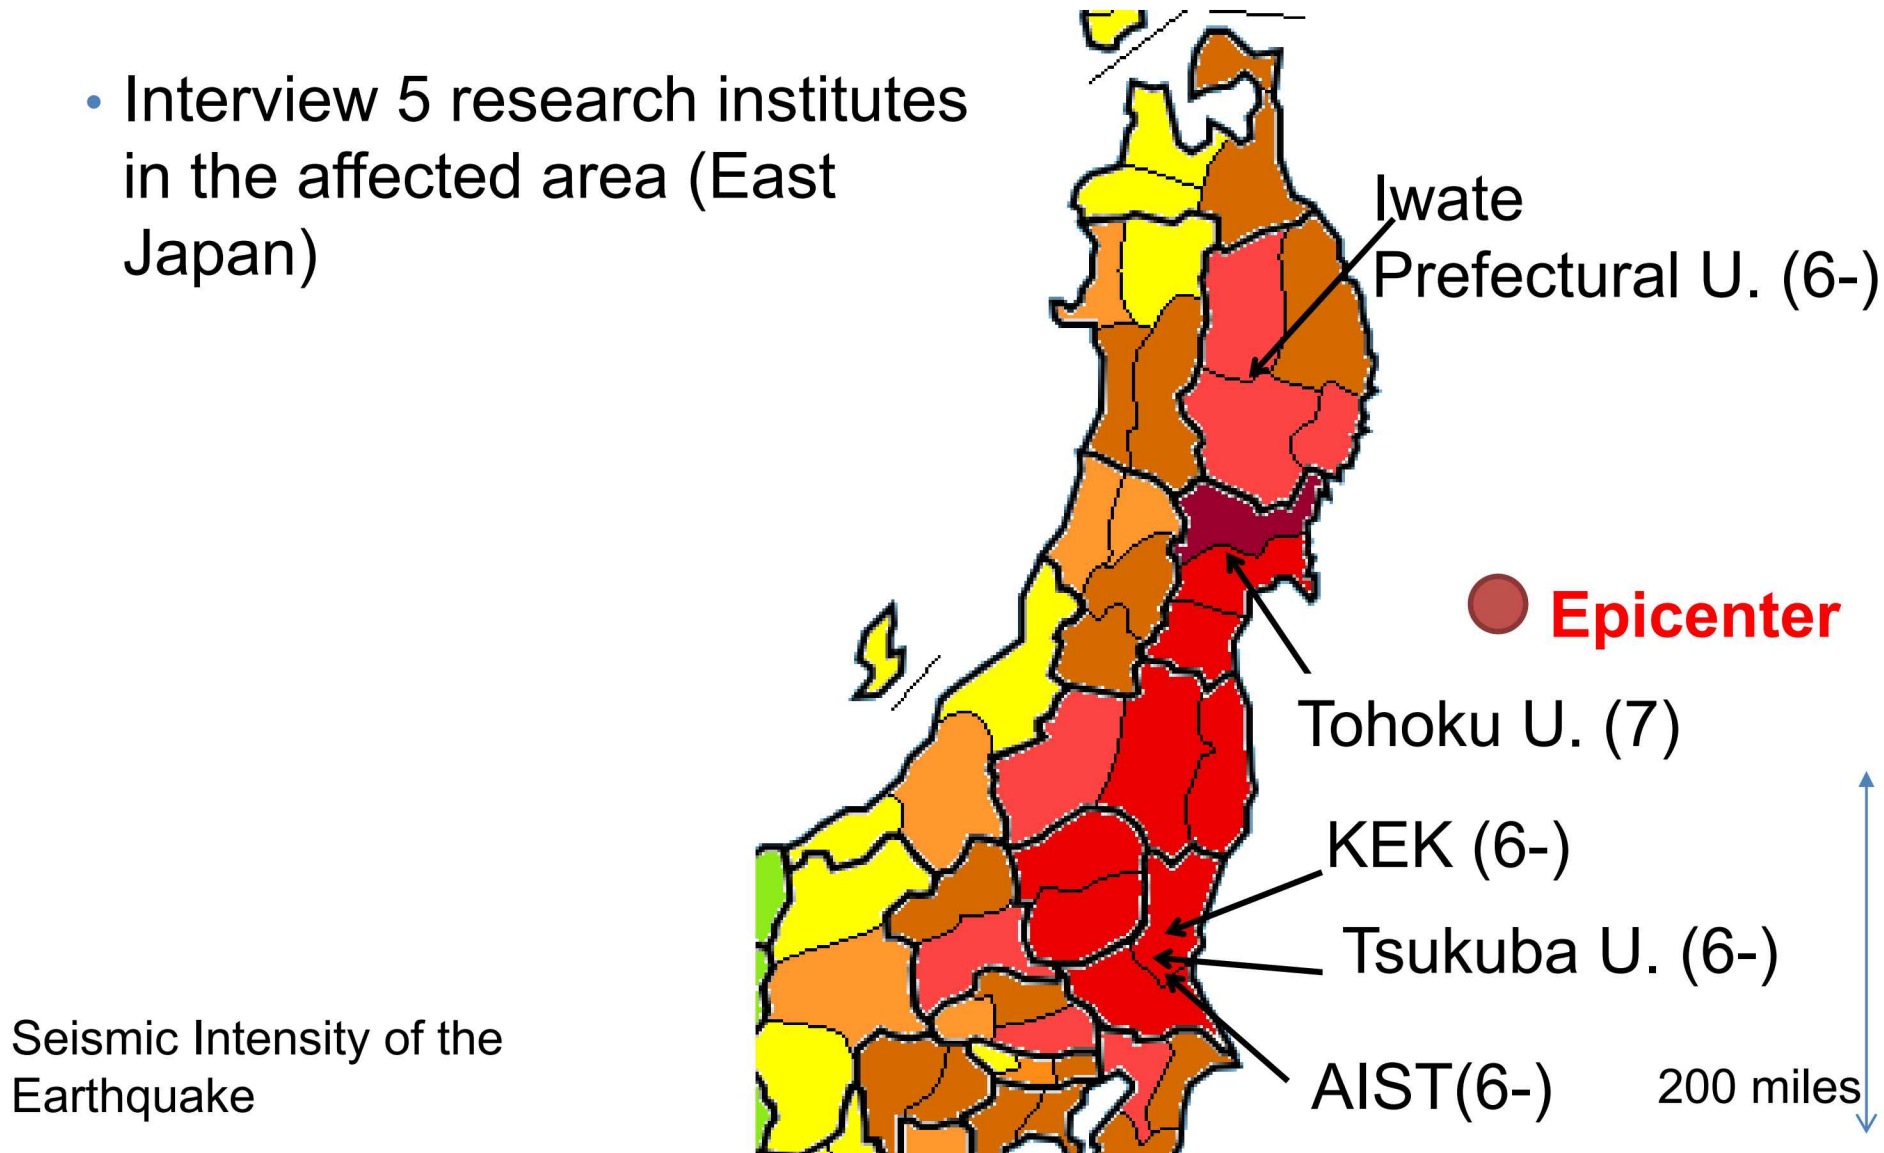
This project studies the effectiveness of movable *virtualized datacenters* in keeping IT services alive during and after a disaster by investigating the joint usage of VM migration (live or using checkpoints), virtual networking, and shared/replicated storage for VM images.

The efforts focused on the following thrusts:

1. Analysis of data and events associated with damaged IT services due to the Great East-Japan Earthquake.
2. Scalability studies of *wide-area VM live-migration*
3. Scalability studies of *wide-area VM backup and check-pointing*
4. An architecture to deploy IT infrastructures in virtualized and distributed datacenters that is resilient to partial physical infrastructure failures

What happened in datacenters?

- Interview 5 research institutes in the affected area (East Japan)

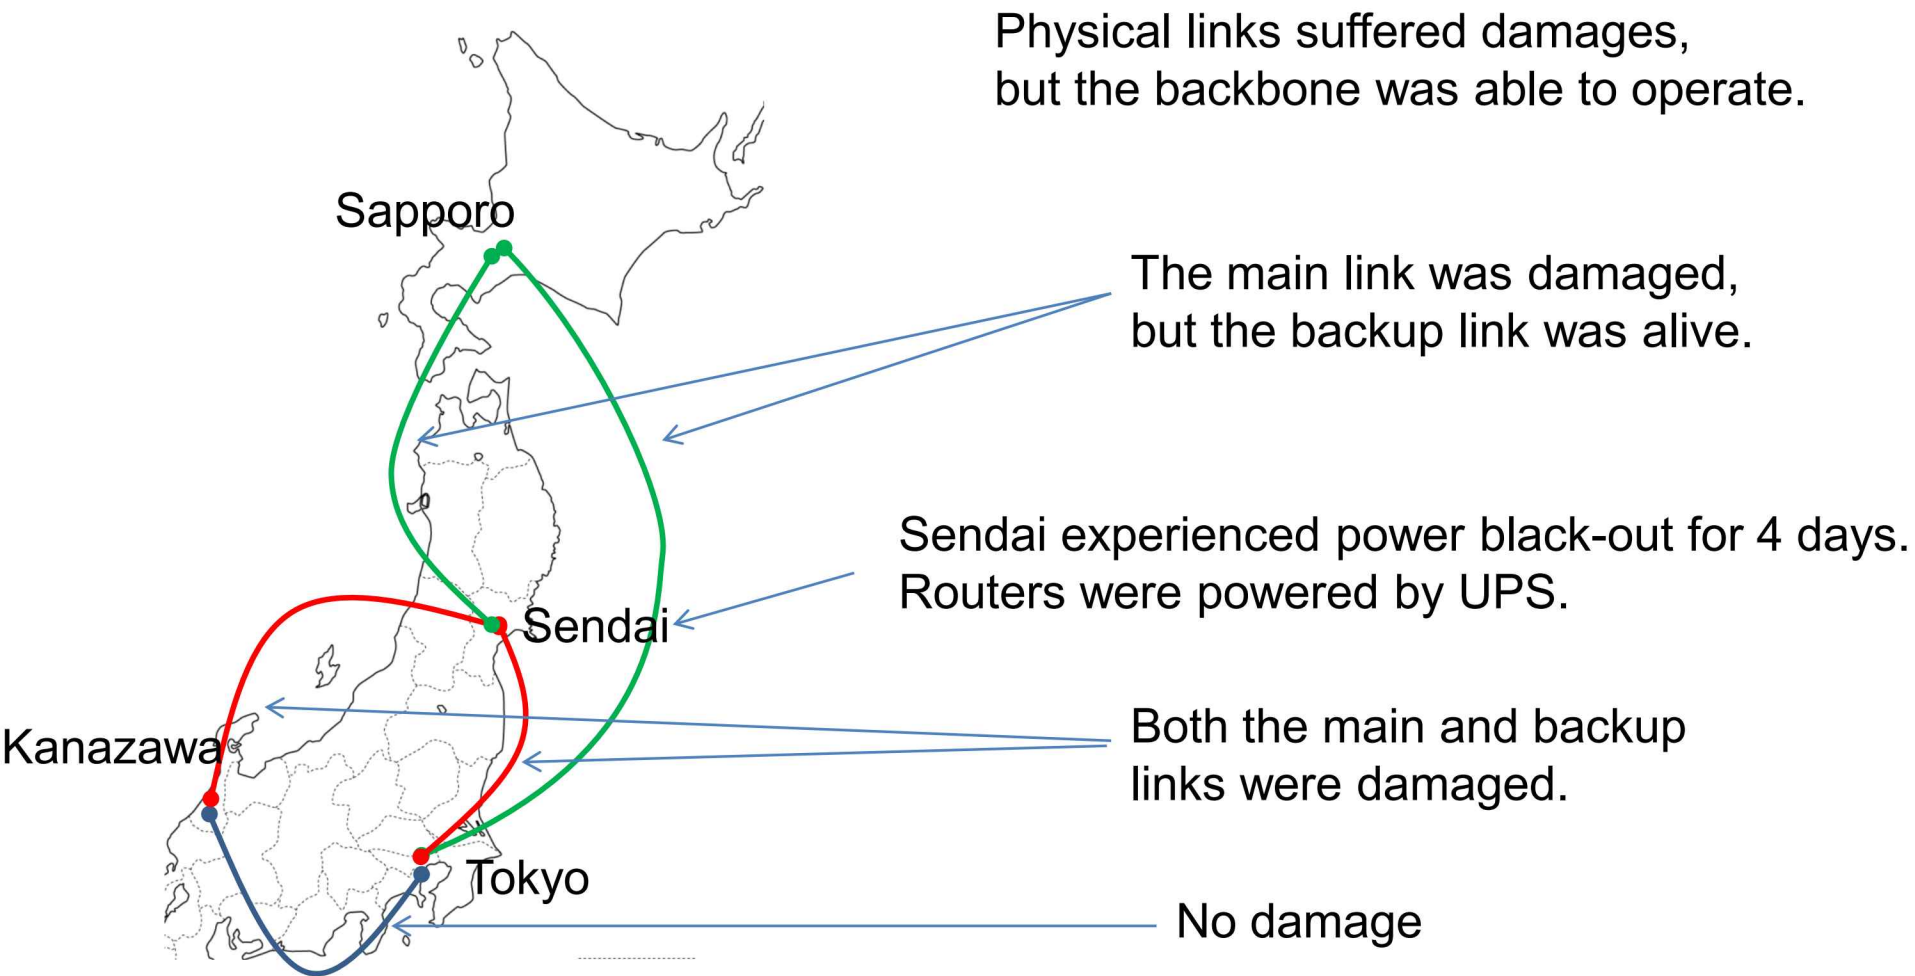


Damages to Datacenters

	<i>Distance from the Epicenter</i>	<i>Seismic Intensity</i>	<i>IT equipment damages</i>	<i>Electrical Power</i>	<i>Network Connectivity</i>
Iwate Prefectural University (岩手県立大学)	220 km	6-	none	Power uninterrupted (generators)	Redundant links kept connectivity alive
Tohoku University (東北大学)	150 km	6- to 6+	none	UPS supplied tens of minutes	Lost after 28 minutes, due to SINET shutdown
KEK (高エネルギー加速器研究機構)	310 km	6-	none	UPS supplied tens of minutes	Data not available
Univ. of Tsukuba (筑波大学)	310 km	6-	none	UPS supplied tens of minutes	Lost immediately
AIST (産業技術総合研究所)	310 km	6-	minimal	UPS supplied 15 to 60 minutes	Available for 60 minutes

Electricity from power company became down just after the earthquake, and the blackout continued for 1-4 days. Most servers and the Internet were alive for tens of minutes.

Damage of SINET4 (One of the Major Academic Networks in Japan)



This slide is based on the information reported in NII Today Vol.52 “東日本大震災でもサービスの提供を続けていたSINET4”.

http://www.nii.ac.jp/userdata/results/pr_data/NII_Today/52/p8-9.pdf

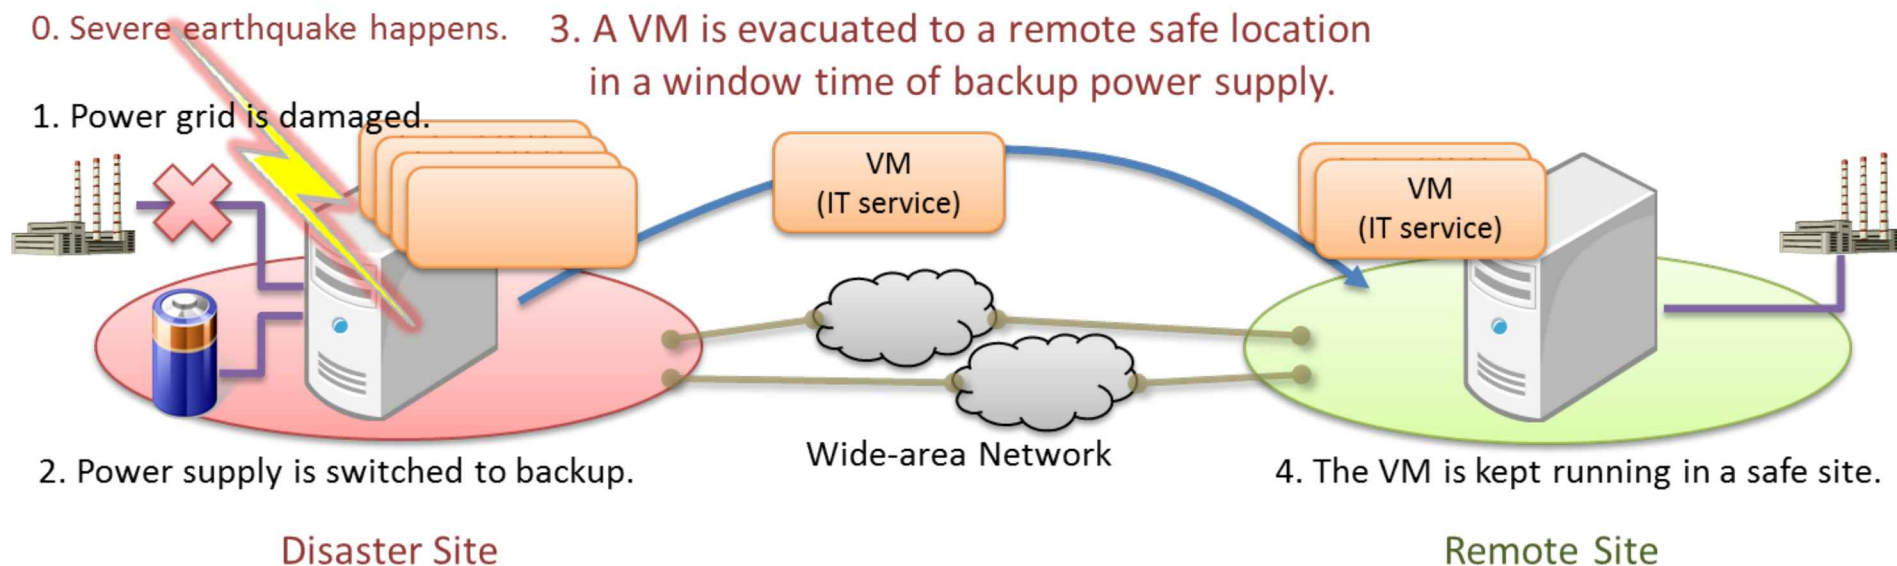
Key Findings

- Our interviews revealed new findings regarding damages of IT infrastructure upon the severe earthquake.
 1. *Most of IT equipment operational during and after the quake,*
 2. *Electrical power available for 30 to 60 minutes,*
 3. *Network connectivity available for 30 to 60 minutes.*
- There is the high possibility that virtualized servers can be evacuated to safe locations upon severe disasters by using modern migration technologies.


On the Use of Virtualization Technologies to Support Uninterrupted IT Services,
IEEE ICC2012 Workshop on Re-think ICT infrastructure designs and operations, Jun 2012

Movable Datacenter Concept

- Evacuate IT services to a safe location in a 60-minutes window of time, while servers and network are operational by power backup.
- Can we use *state-of-art virtualization technologies* for the evacuation?
 - It can relocate a VM to another host transparently.
 - No visible interruption to applications
 - No special program required to run on VM.
 - But, designed for LAN environments.



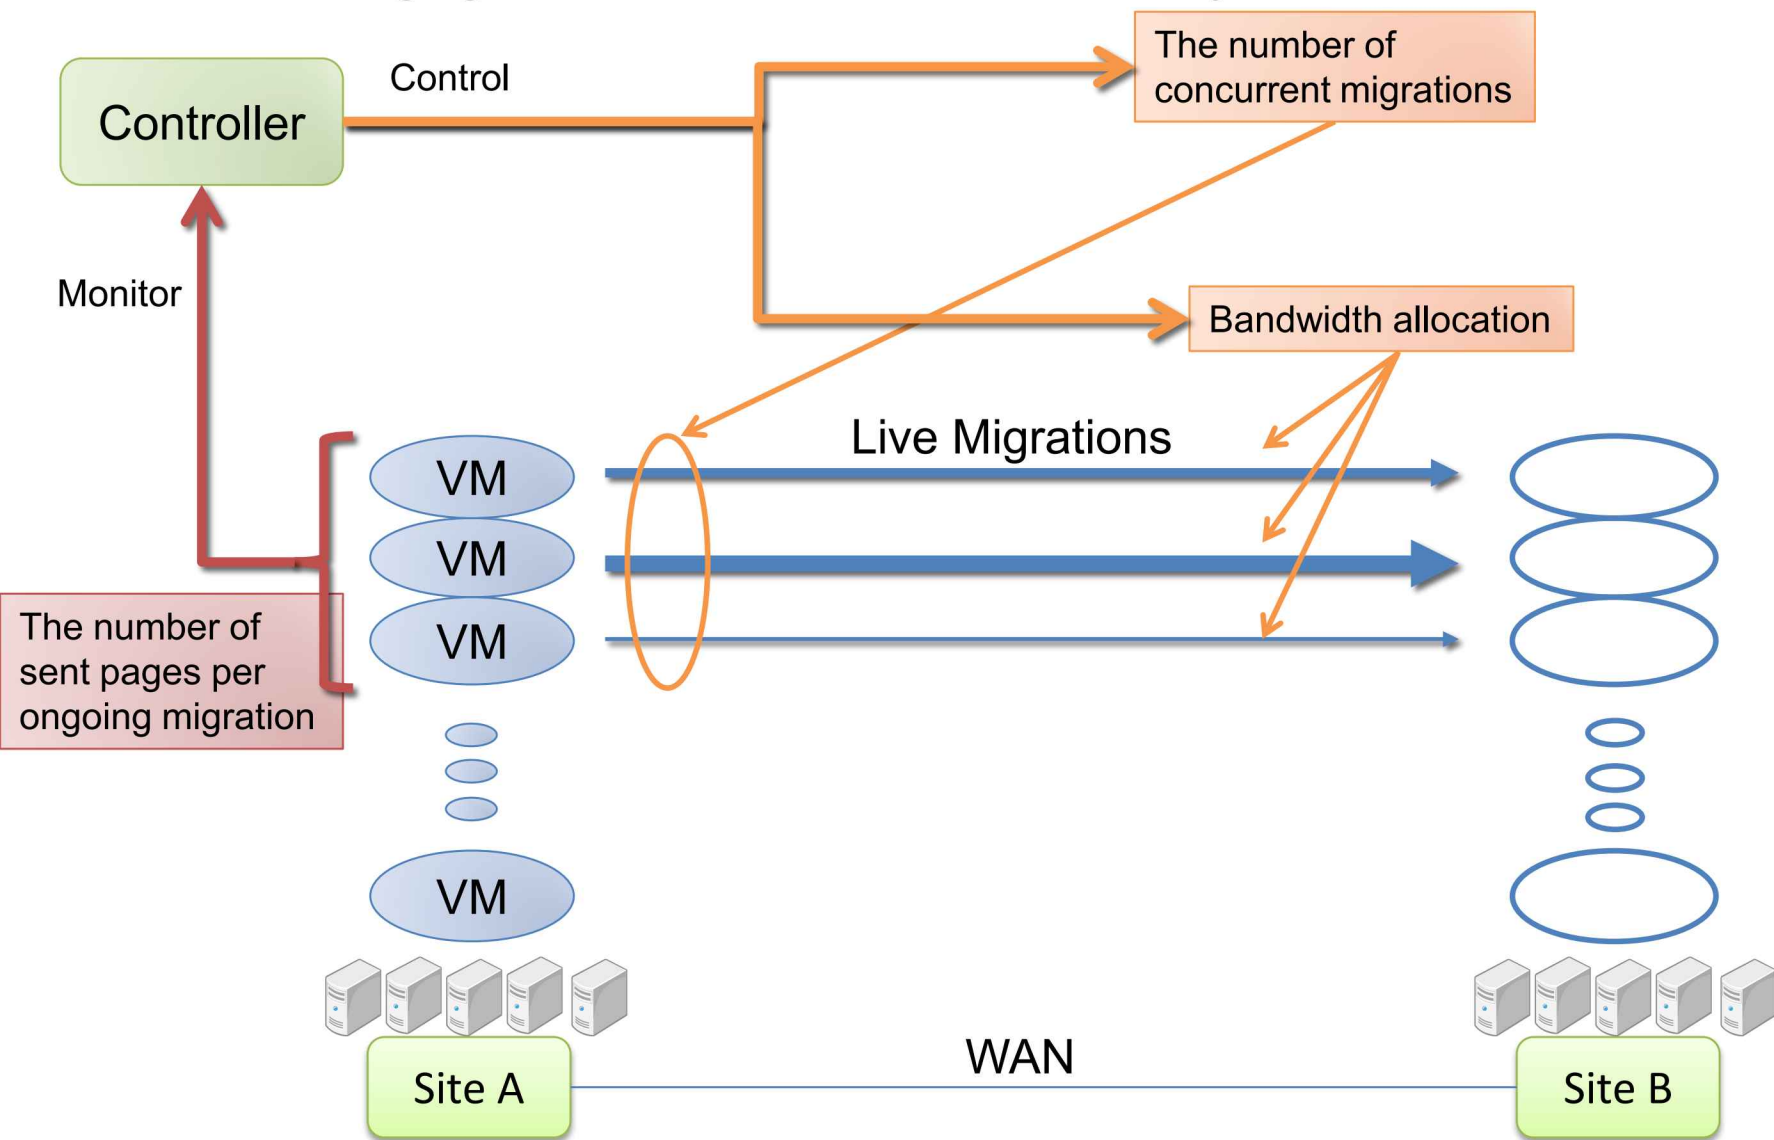
Can state-of-art virtualization technologies safely evacuate IT systems upon a disaster?



- We evaluated virtual machine technologies under a real-world long-distance network over the Pacific.
- We observed
 - A poor network condition adversely affects individual migration time.
 - Parallel live migrations increase evacuation throughput of VMs in WANs, but also increase the risk of evacuation failure.
 - Disaster recovery mechanisms incur performance degradation in normal operations.

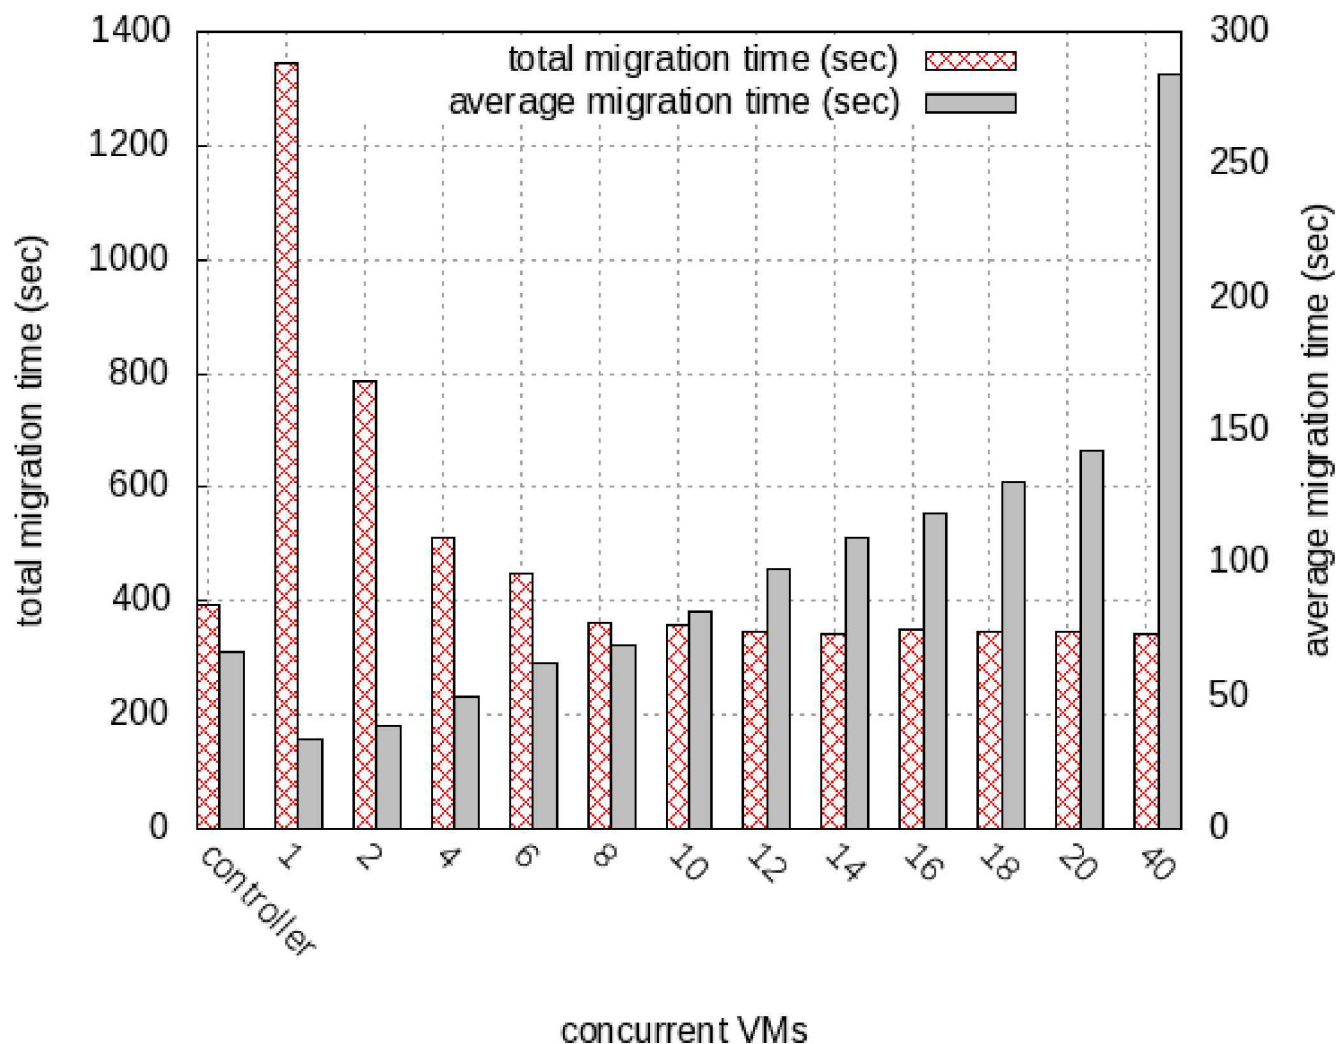
Automatic Feedback Control

Achieve maximum evacuation throughput with minimum per-VM migration time for changing network conditions and VM activity.




Advantages of Feedback Control

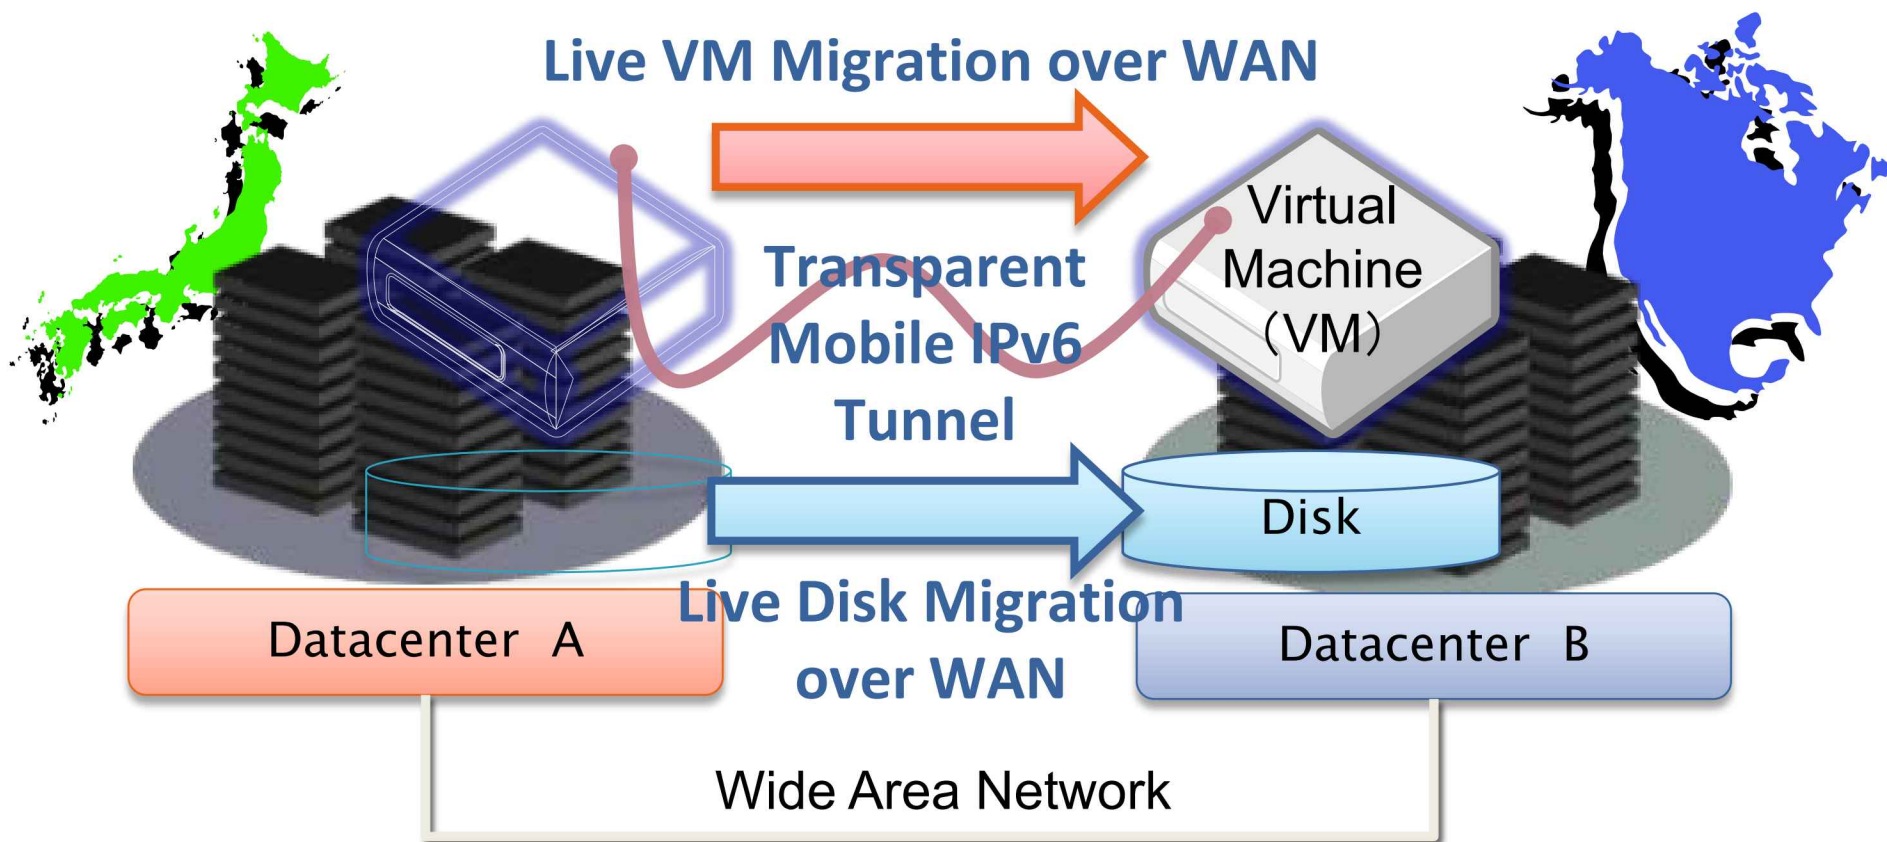
Evacuate 40VMs to a safe location over an unstable network



Controller in action



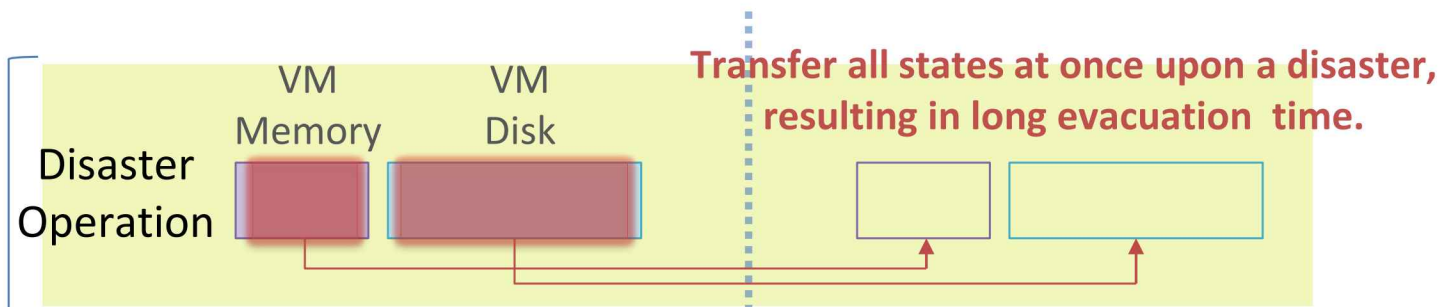
Inter-datacenter Migration Protocol (1)



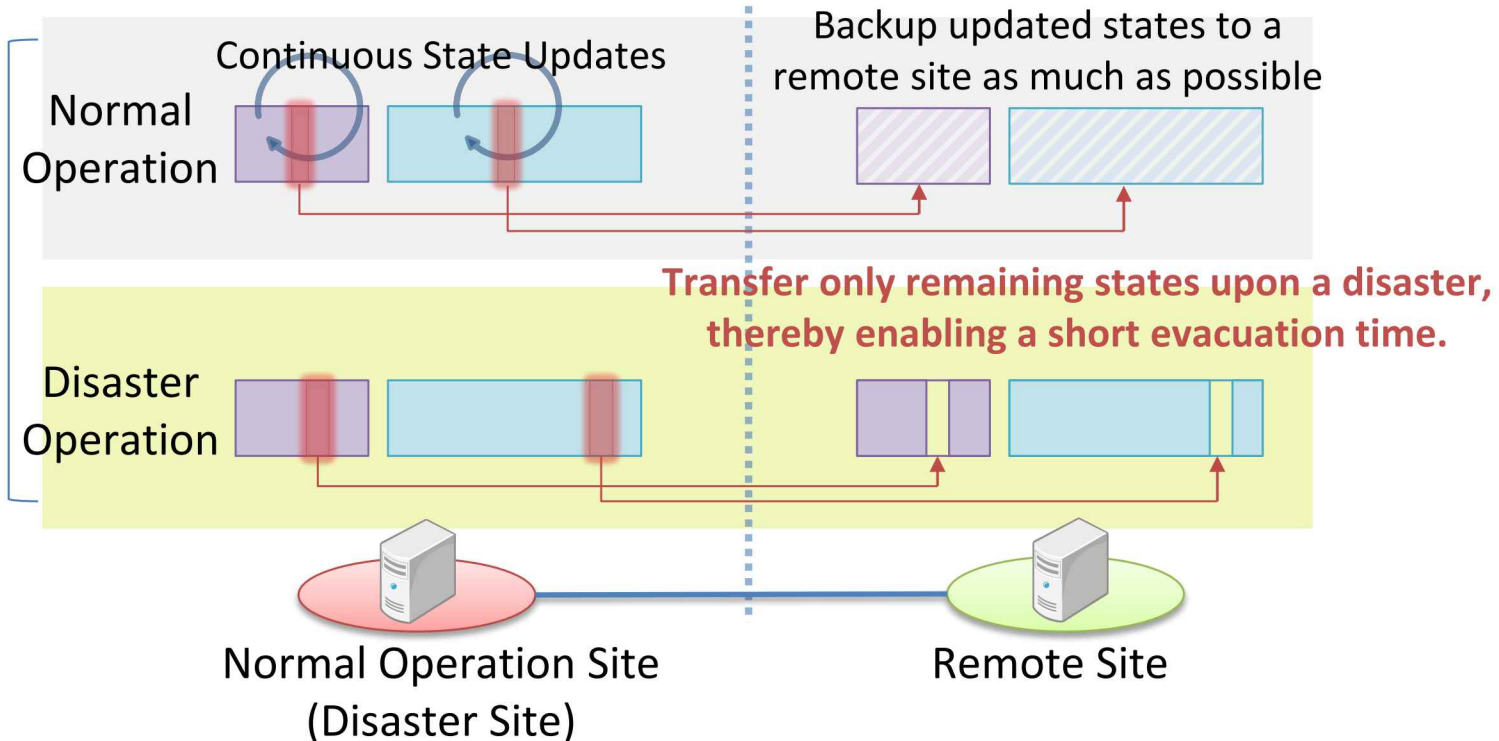
A WAN-optimized Live Storage Migration Mechanism Toward Virtual Machine Evacuation Upon Severe Disasters, In submission to IEICE Transactions on Information and Systems, 2013

Inter-datacenter Migration Protocol (2)

Traditional Migration Protocol



Inter-datacenter Migration Protocol

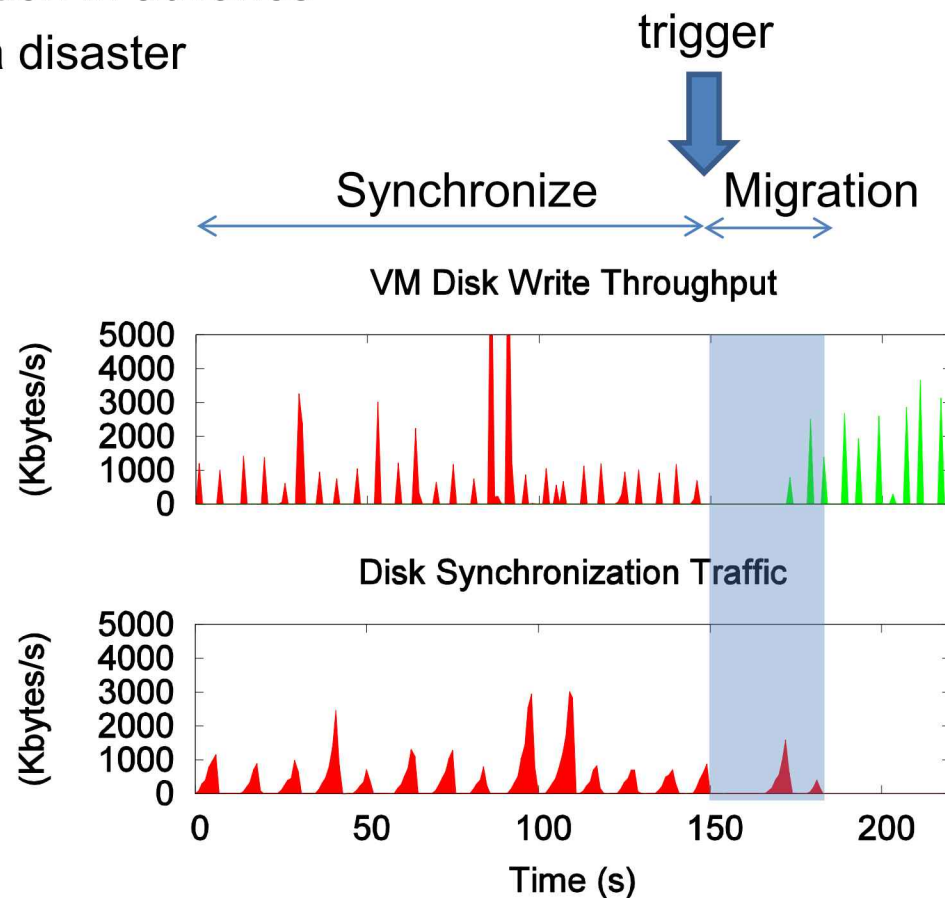


Inter-datacenter Migration Protocol (3)

- Our preliminary experiments confirmed that the mechanism successfully shortens individual live migration time
 - Synchronize VM states to destination in advance
 - Copy the rest of VM states upon a disaster

Prototype

- The whole VM states including its virtual disk *migrated over the Pacific Ocean just in 30 seconds.*
- 512 RAM, 4GB virtual disk



Publications

- On the Use of Virtualization Technologies to Support Uninterrupted IT Services
 - M. Tsugawa, R. Figueiredo, J. Fortes, T. Hirofuchi, H. Nakada, and R. Takano
 - IEEE ICC2012 Workshop on Re-think ICT infrastructure designs and operations, June 2012.
- Lessons Learnt from a Preliminary Prototype of a Best-Effort Pre-synchronization Mechanism for Wide-Area Live Migration of Virtual Machines (Work-in-Progress Report)
 - T. Hirofuchi, M. Tsugawa, H. Nakada, S. Itoh, and S. Sekiguchi
 - Information Processing Society of Japan SIG Technical Report, May 2012
- Reducing the Migration Times of Multiple VMs on WANs
 - T. S. Kang, M. Tsugawa, T. Hirofuchi, and J. Fortes
 - ACM Student Research Competition Poster, SC'12, November 2012
- A WAN-optimized Live Storage Migration Mechanism Toward Virtual Machine Evacuation Upon Severe Disasters
 - T. Hirofuchi, M. Tsugawa, H. Nakada, T. Kudo, and I. Satoshi
 - In submission to IEICE Transactions on Information and Systems, 2013

Conclusion

- Lessons learned from the Great East Japan Earthquake
 - In datacenters, physical damages to servers were minimum.
 - Servers were operational for tens of minutes by power backup.
 - Internet connectivity was available if switches were operational.
- Datacenter evacuation on extreme events
 - Evacuate IT services to a safe location in a limited time window
 - Study live VM migration for WAN environments
- Further research and development needed
 - Improve VM migration performance
 - Intelligent and efficient use of resources
 - Electrical power resiliency – improve battery/generator backup
 - Network infrastructure resiliency