Seamless "Live Virtual Machine Migration" by mitigating “shared storage” constraint

Sangeeth Keeriyadath
Prasanth Jose

IBM

Agenda

- Live Virtual Machine migration
  - Overview
  - Reasons
  - Requirements
- Shared Storage requirement
  - Musing – Shared or Local
  - Challenges for “shared storage” setup
- Methods to allay “shared storage” constraint
  - Across storage protocol boundaries
  - Remote storage access
  - Efficient storage replication

Live Virtual Machine(VM) Migration

Live Virtual Machine(VM) Migration is the movement of a running VM from one physical machine to another without disrupting the operation of the Operating System and Applications running in it.

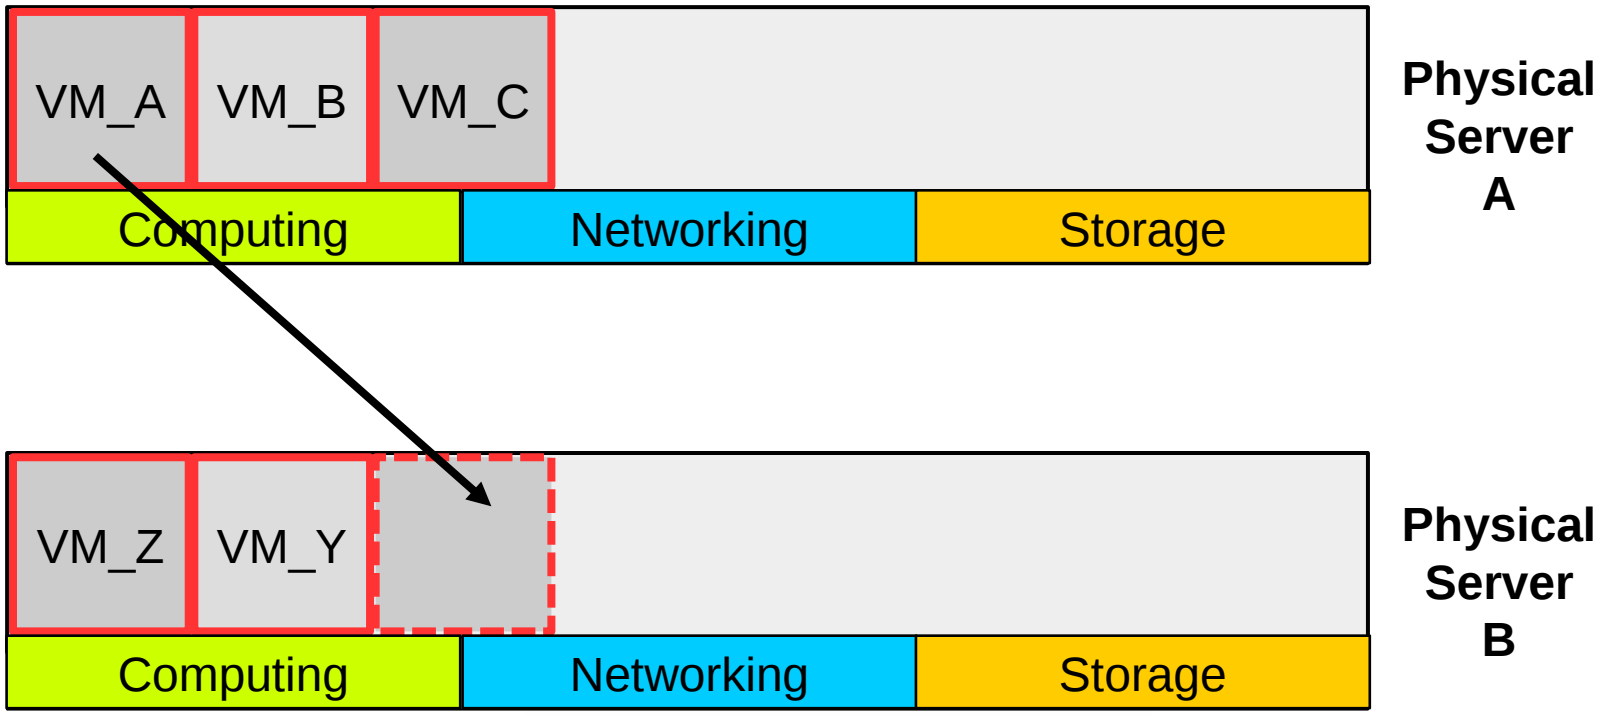

Reasons for VM migration

1. Load distribution across servers

- Migrate to a newer server in the data center
- Migrate to a server with more available resources

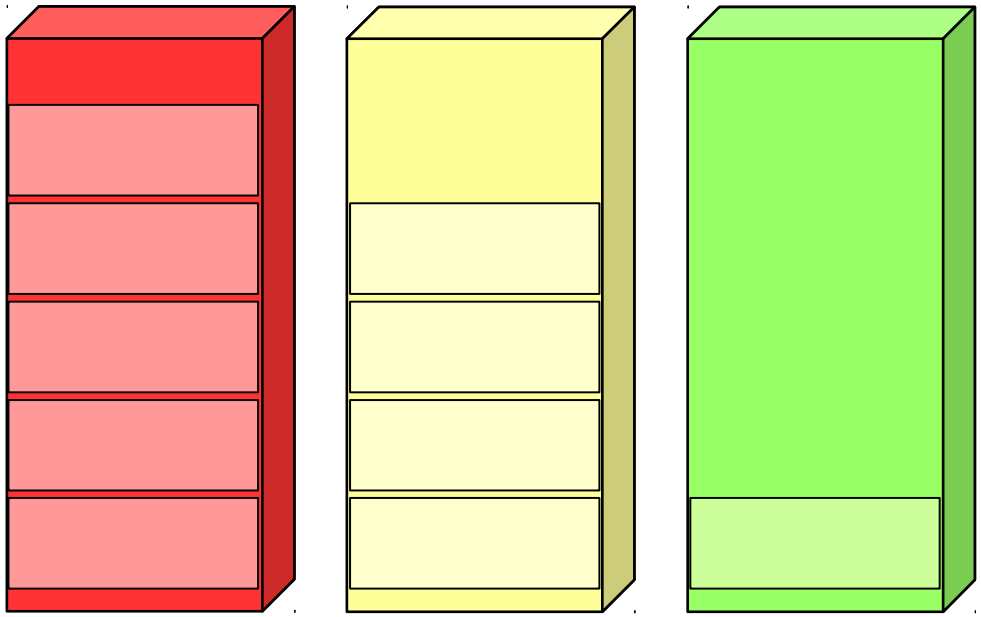

Reasons for VM migration

2. Server Consolidation

- Increase server utilization and decrease operating costs
- Consolidate all the idle or non-critical virtual machines to a different server

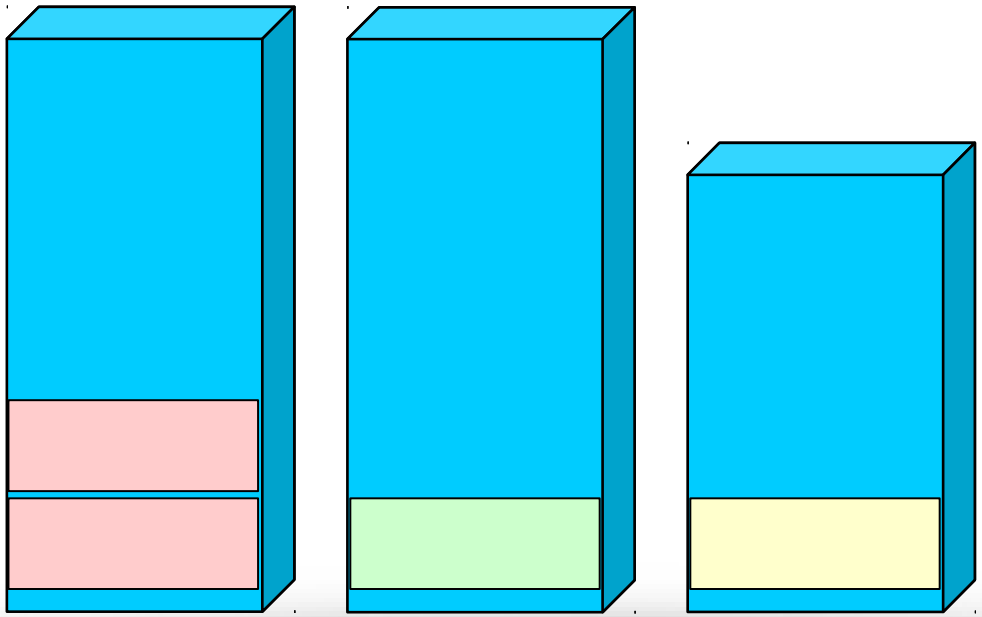

Reasons for VM migration

3. Server maintenance

- Migrate to an alternate server without workload downtime. Source server can be taken down for maintenance of hardware components and/or firmware etc.

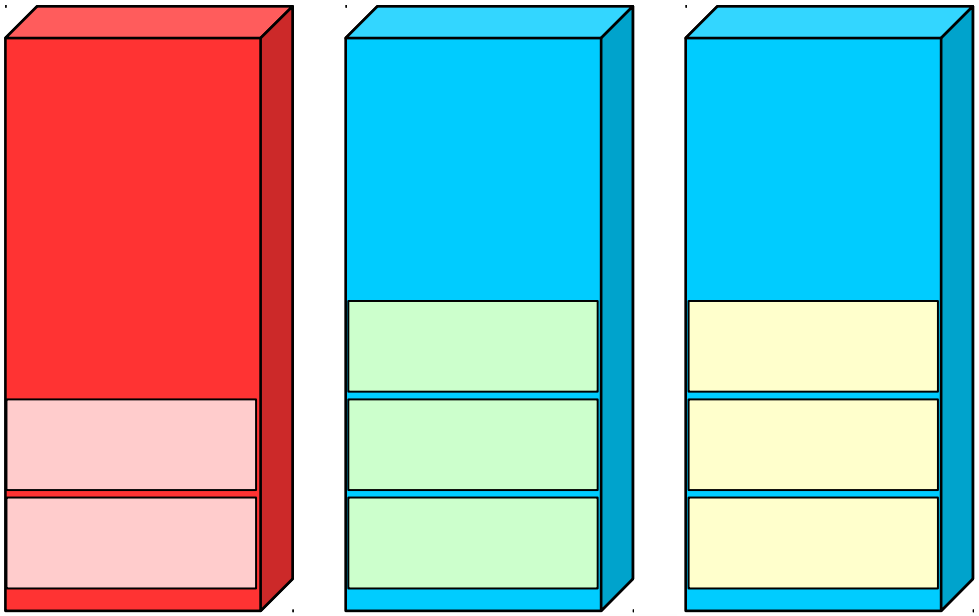

Requirements for the destination server to successfully host an incoming VM

- **Computing** – Sufficient processor/memory
- **Networking** – Connectivity to same ethernet Network
- **Storage** – Connectivity to ***same(shared)*** storage devices

Musing – Shared or Local storage (1)

- ✓ Scales up easily
- ✓ Shared across multiple servers - suits data center growth
- ✓ Meets Performance and Capacity needs - Hybrid cached model
- ✓ High Availability - Tolerant to Host / Controller failure
- ✓ Flexibility of Unified storage (Block and NAS)

- ✘ Switching and cabling layout could be intimidating
- ✘ Speed and distance requirements drive the cost of inter-links
- ✘ Changes could make it horifically slow
- ✘ Havoc, if misconfigured. Business Risk!

Musing – Shared or Local storage (2)

- ✓ Cheaper and Faster (period!)
- ✓ Secure
- ✓ Simplicity drives stability

- ✘ Doesn't scale well
- ✘ Unused storage can't move to another server
- ✘ Backup and Archiving complexities

Challenges with having shared storage setup for VM migration

- Data center growth / (re)design
- Cost/Complexity overhead of maintaining connectivity of same storage across multiple servers
- Heterogeneous connectivity/devices/Switching infrastructure
  - DAS (Internal Storage)
  - FC (Storage Area Network)
  - iSCSI (SCSI over standard ethernet)
  - FCoE (Converged network – preserving FC constructs)

Reference architecture for the proposed solutions : IBM® PowerVM®

IBM® PowerVM® provides the industrial-strength virtualization solution for IBM Power Systems™ servers and blades that run IBM AIX®, IBM i and Linux workloads.

PowerVM technologies referenced

- **Live Partition Mobility (LPM)** : Technology that migrates a running VM from one Power Server to another without any application downtime.
- **Virtual I/O Server (VIOS)**
  - ➔ **Virtual networking** : VIOS supports software switch systems that use shared ethernet adapter technology.
  - ➔ **Virtual SCSI** : VIOS does the storage virtualization. It performs SCSI emulation and acts as SCSI target.
  - ➔ **NPIV** : VIOS facilities storage adapter sharing and there is no device level abstraction or emulation.

Methods to alleviate “shared storage” constraint

- **Method A** : VM migration across storage protocol boundaries
- **Method B** : VM migration with slower access to remote storage
- **Method C** : VM migration with storage replication

Method A : VM migration across storage protocol boundaries

Flexibility of using available “storage transport” on the source and destination machines based on their connectivity to storage.

Benefits of newer Shared Storage boxes :

- Support different storage transport protocols.
- Flexibility for the same storage Logical Unit (Disk) to be mapped to both FC port and iSCSI (or Ethernet) port.

Example :

Connectivity of the same storage Logical Unit to source machine is via FC and to destination machine is via iSCSI. With some modifications to VM migration “device identification and configuration code” we could successfully migrate a VM from one machine to another.

(A) Figure : Across storage protocol boundaries

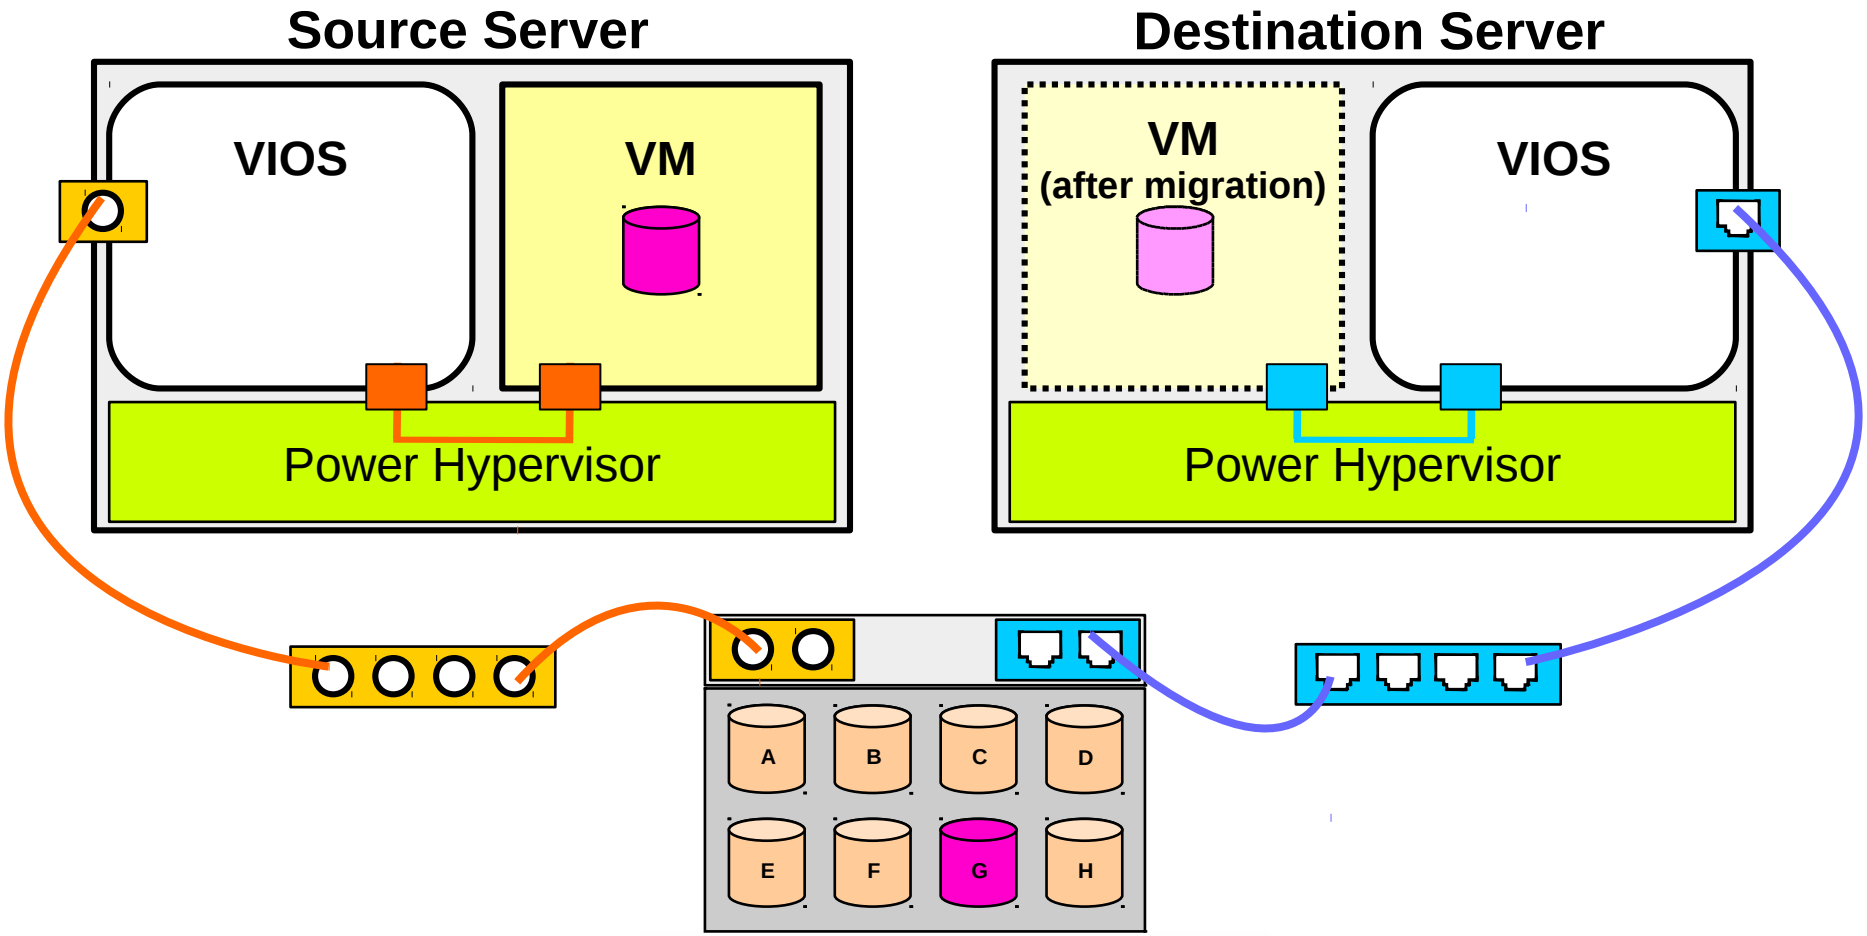

Method B : VM migration with slower access to remote storage

Workloads with intense "**computing resource requirements**" that meet one/more of these criteria :

- Run at a certain time of day/week/month (e.g. Scheduled batch jobs)
- Are processor/memory intensive (e.g. In-memory database)
- Expected to complete as soon as possible (e.g. Financial reports. Sooner the better!)

Solution : Distribute the workload to different servers across the data center based on resource availability.

Used to advantage : Source server need not be powered down.

Example :

Source server exports the storage disks mapped to the client as iSCSI target devices and destination server accesses them as iSCSI initiator.

(B) Figure : For computing resource requirements

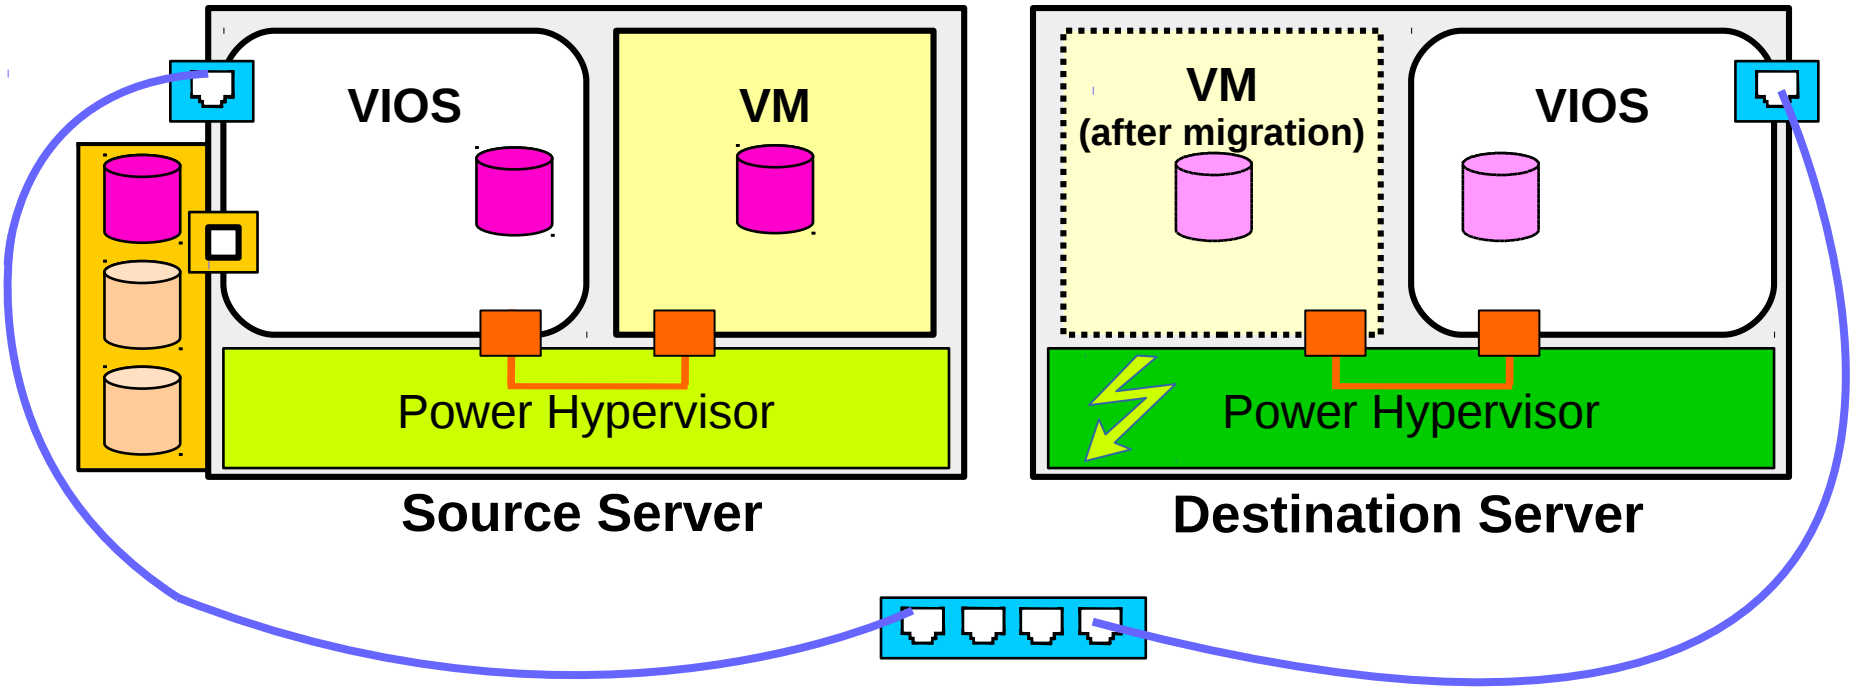

Method C : VM migration with storage replication

"VM migration" with the following characteristics :

- Source server needs to be brought down for maintenance
- Frequent movement of VMs, back and forth the same servers

Challenge : Storage of the VMs are on the local disks of the server and won't be accessible after the source sever is shut down.

Solution : Replicate the VM's storage to destination server

Making it efficient **: Replication progress counter**

: Block modification bitmap

(C) Figure : Storage replication

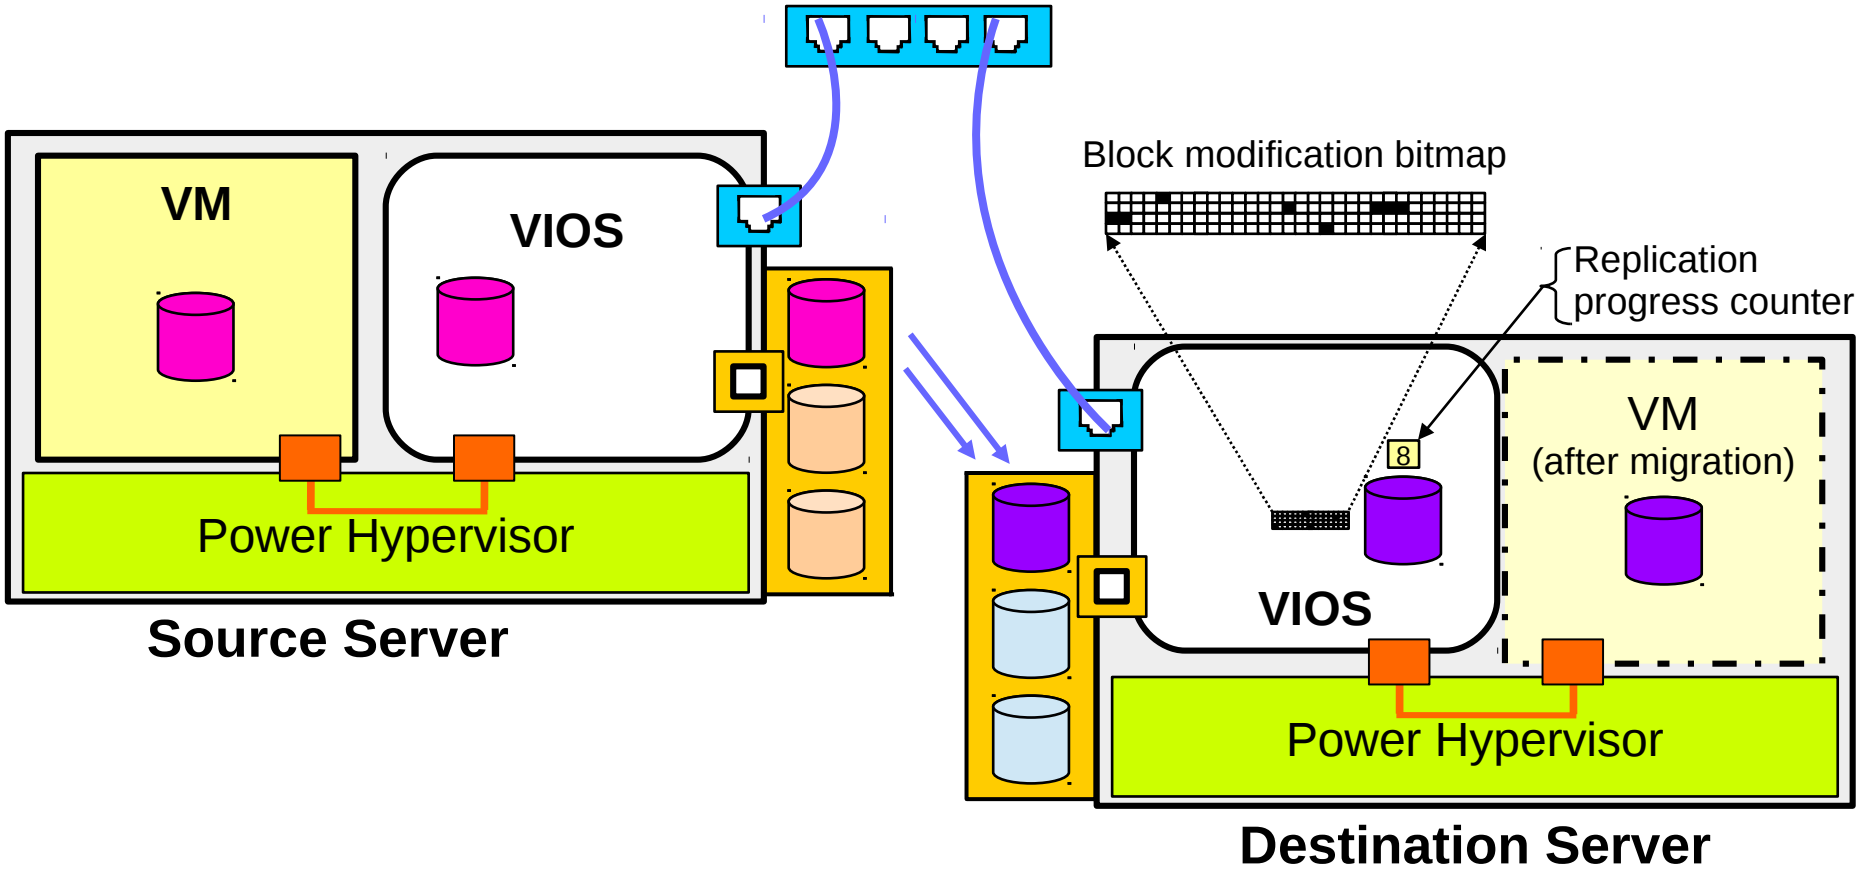

(C) Algorithm for efficient storage replication

- Source server exports the storage disks mapped to the client as **iSCSI target devices**.
- Storage **replication begins after** the VM migrates.
- Replication is performed by destination server which copies fixed chunks of storage blocks in a **sequential and contiguous** manner. Updating the progress in a "replication_progress_counter".
- **Write requests** of the migrated VM are directed to the storage of destination server. A "block_modification_bitmap" of VM's storage blocks is maintained to indicate the blocks written on destination server.

(C) Algorithm for efficient storage replication (continued)

- **Read requests** of the VM(before storage mirroring is complete) are checked against the “replication_progress_counter” and if required, against “block_modification_bitmap”.

- After storage replication; all read and write requests are performed on the destination storage. The “block_modification_bitmap” is **updated continuosly** to indicate the storage blocks modified by VM on destination server.

- “block_modification_bitmap” is used during the migration of VM back to the source server; to **efficiently replicate only the modified blocks**.

(C) VM's I/O flow during "storage replication"

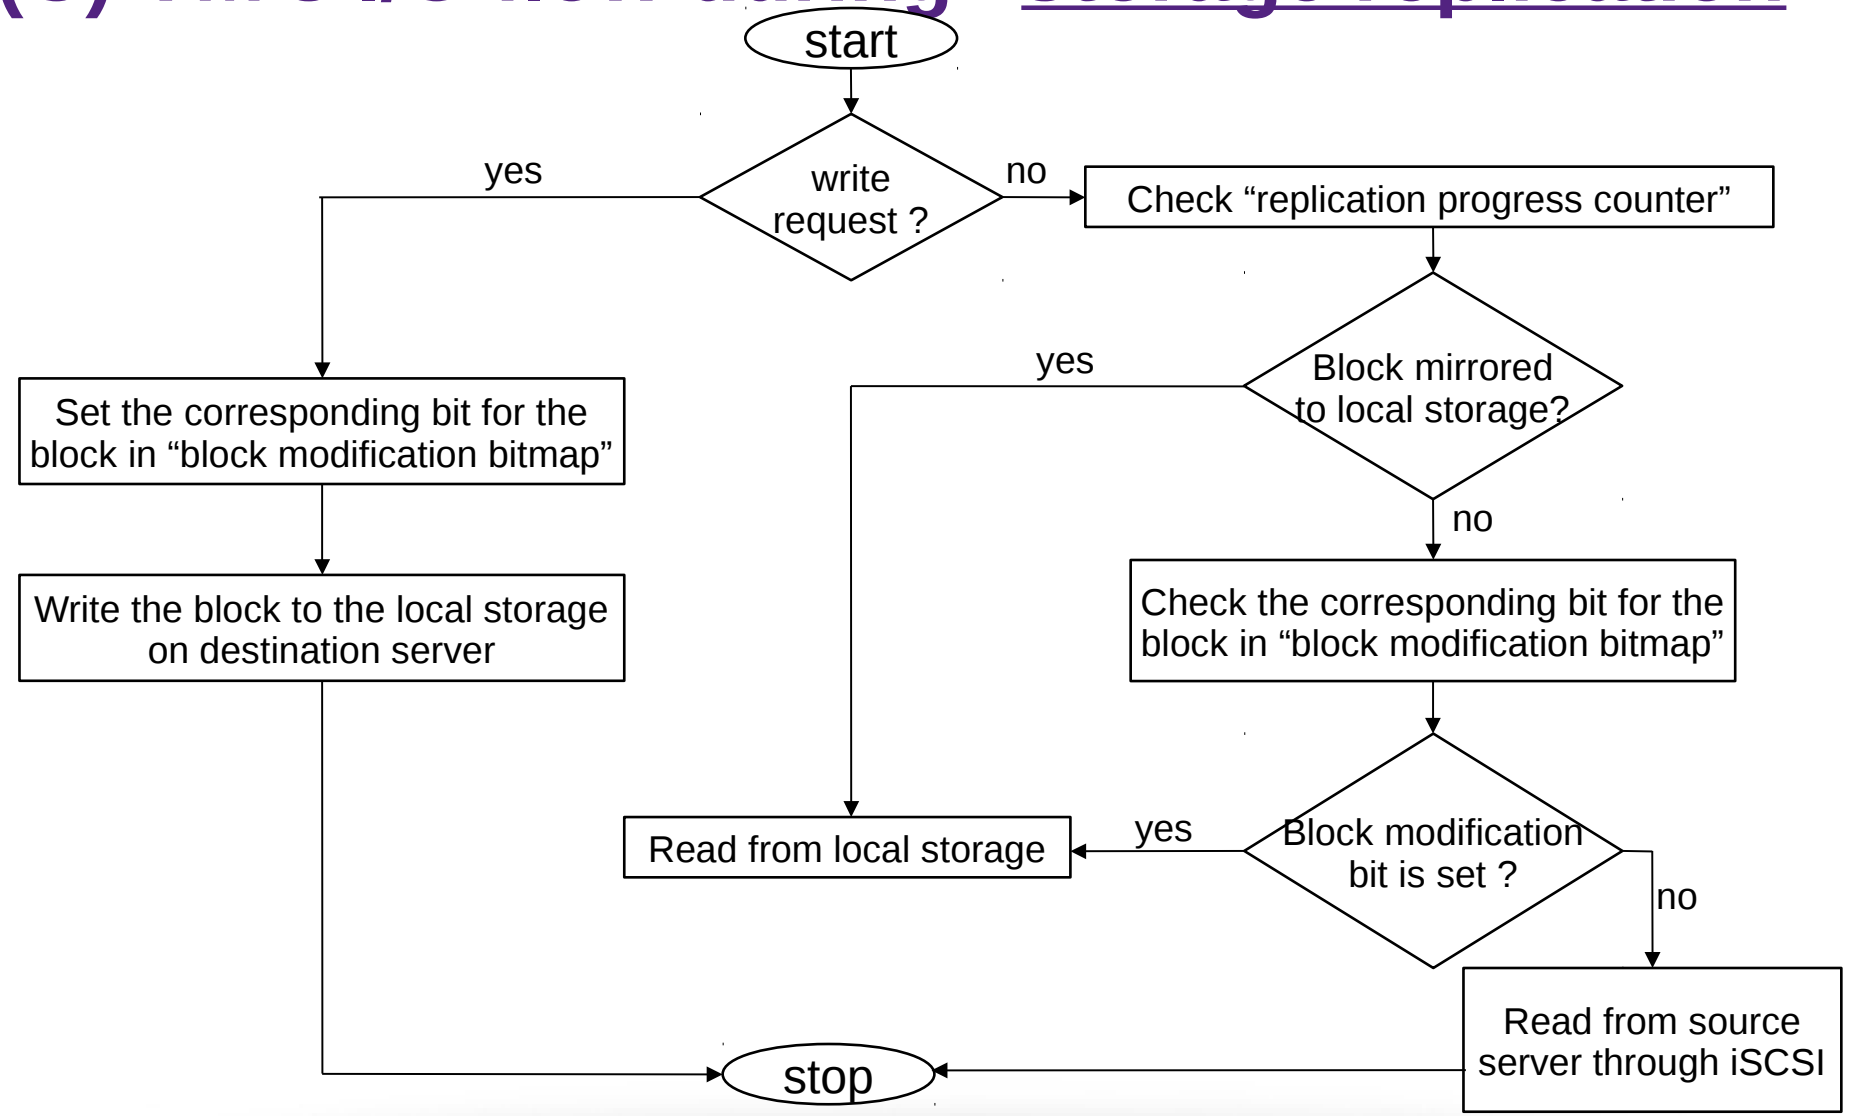

(C) To and Fro "subsequent migration" flow

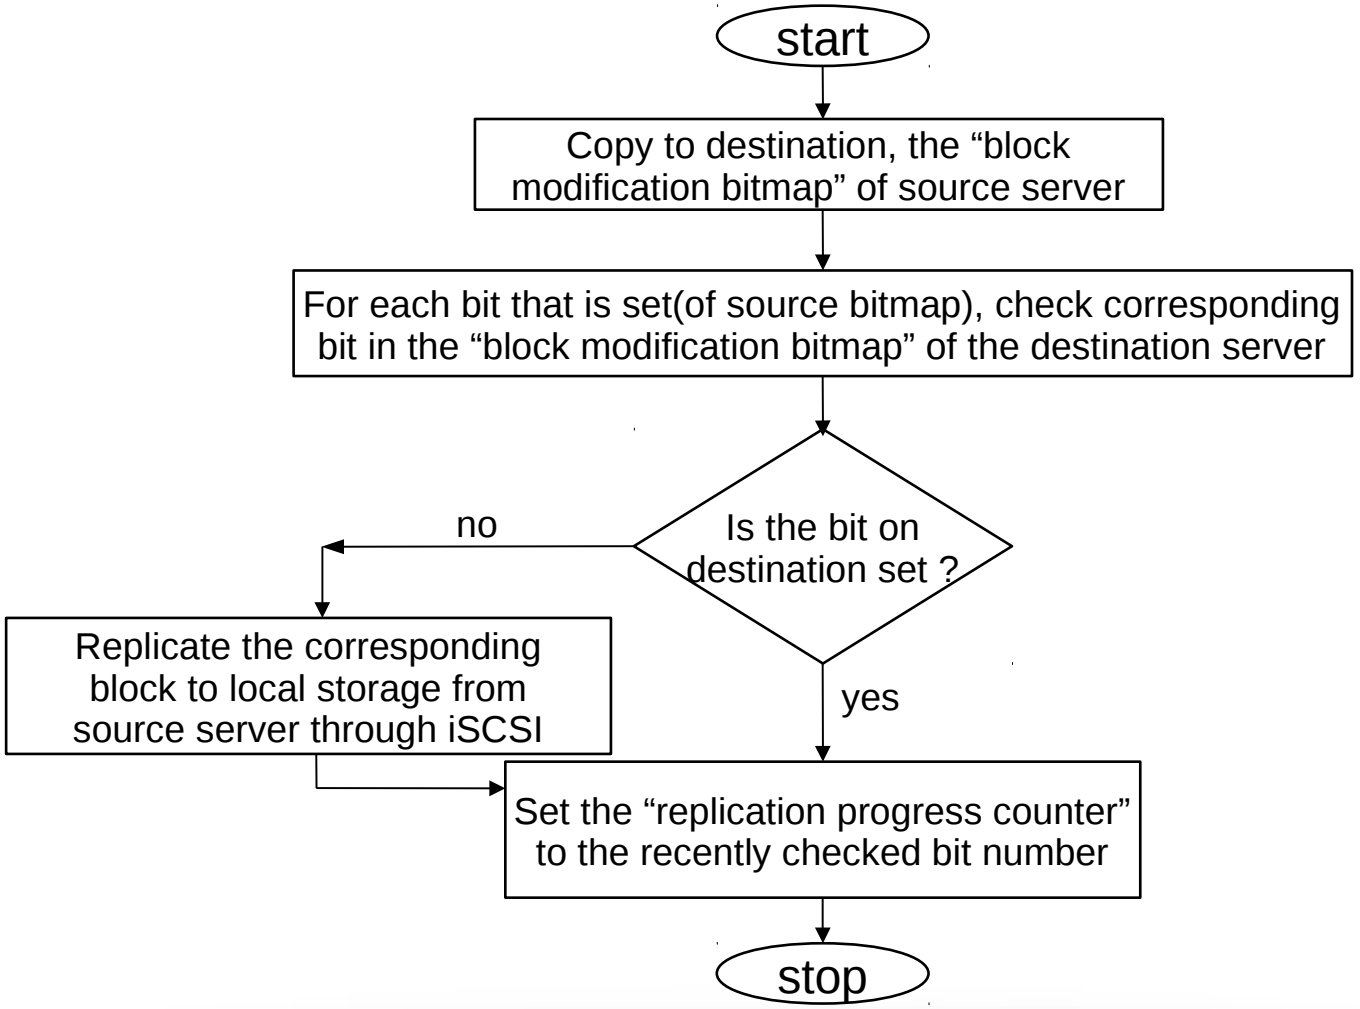

Conclusion

- Improving the **availability** of businesses, with **increased flexibility** and **reduced costs** is a continuously evolving piece of technology and this presentation is our contribution to it.
- We believe to have proposed **key valuable enhancements** to the already existing methods to mitigate the need for shared storage for VM migration.

References

IBM PowerVM Virtualization Managing and Monitoring :
http://www.redbooks.ibm.com/abstracts/sg247590.html?Open

Virtual I/O Server :
http://www-01.ibm.com/support/knowledgecenter/POWER8/p8hb1/p8hb1_kickoff.htm?cp=POWER8%2F1-6-1-3-5

VMWare vSphere migration with vMotion in Environments without shared storage :
http://pubs.vmware.com/vsphere-51/index.jsp?topic=%2Fcom.vmware.vsphere.vcenterhost.doc%2FGUID-561681D9-6511-44DF-B169-F20E6CA94944.html

Openstack VM live migration :
https://kimizhang.wordpress.com/2013/08/26/openstack-vm-live-migration/

Shared-Nothing VM Live Migration with Windows Server 2012 Hyper-V :
http://windowsitpro.com/windows-server-2012/shared-nothing-vm-live-migration-windows-server-2012-hyper-v

World Wide Web !

Thank You