



Education

Virtualization Practices: Providing a Complete Virtual Solution in a Box

Jyh-shing Chen, NetApp

Author: Jyh-shing Chen, NetApp

- The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

➤ Virtualization Practices: Providing a Complete Virtual Solution in a Box

- ◆ This session will appeal to Data Center Managers, Development Managers, and those that are seeking a fundamental understanding of deploying virtualized solutions in a box. The session will delve into the benefits of providing virtualized solutions and its value. Not limited to developers, the audience will receive the concepts in deploying virtualized solutions in a box and how they can perform tuning and troubleshooting in this type of environment.

- Virtual solution in a box comes in various form factors.



(From: commons.wikipedia.org)

- This tutorial will be focusing on virtualization within a single server/computer.

➤ Quicker deployment

- ◆ Reduce rack and stack efforts
- ◆ Easier resource management

➤ Better resource utilization

- ◆ Consolidate physical resources
- ◆ Reduce data center footprint and operating costs

➤ Provide hardware abstraction

- ◆ Run on virtual hardware platform
- ◆ Allow easy migration to different physical hardware platform

➤ Easier troubleshooting

- ◆ Adjust virtual resources to tune the solution
- ◆ Deploy pre-configured troubleshooting virtual machine

➤ Solution Components

- ◆ Compute
- ◆ Network
- ◆ Storage

➤ Solution Architectures

- ◆ Shared infrastructure
- ◆ Isolated environment
- ◆ Hybrid environment

➤ Resource management

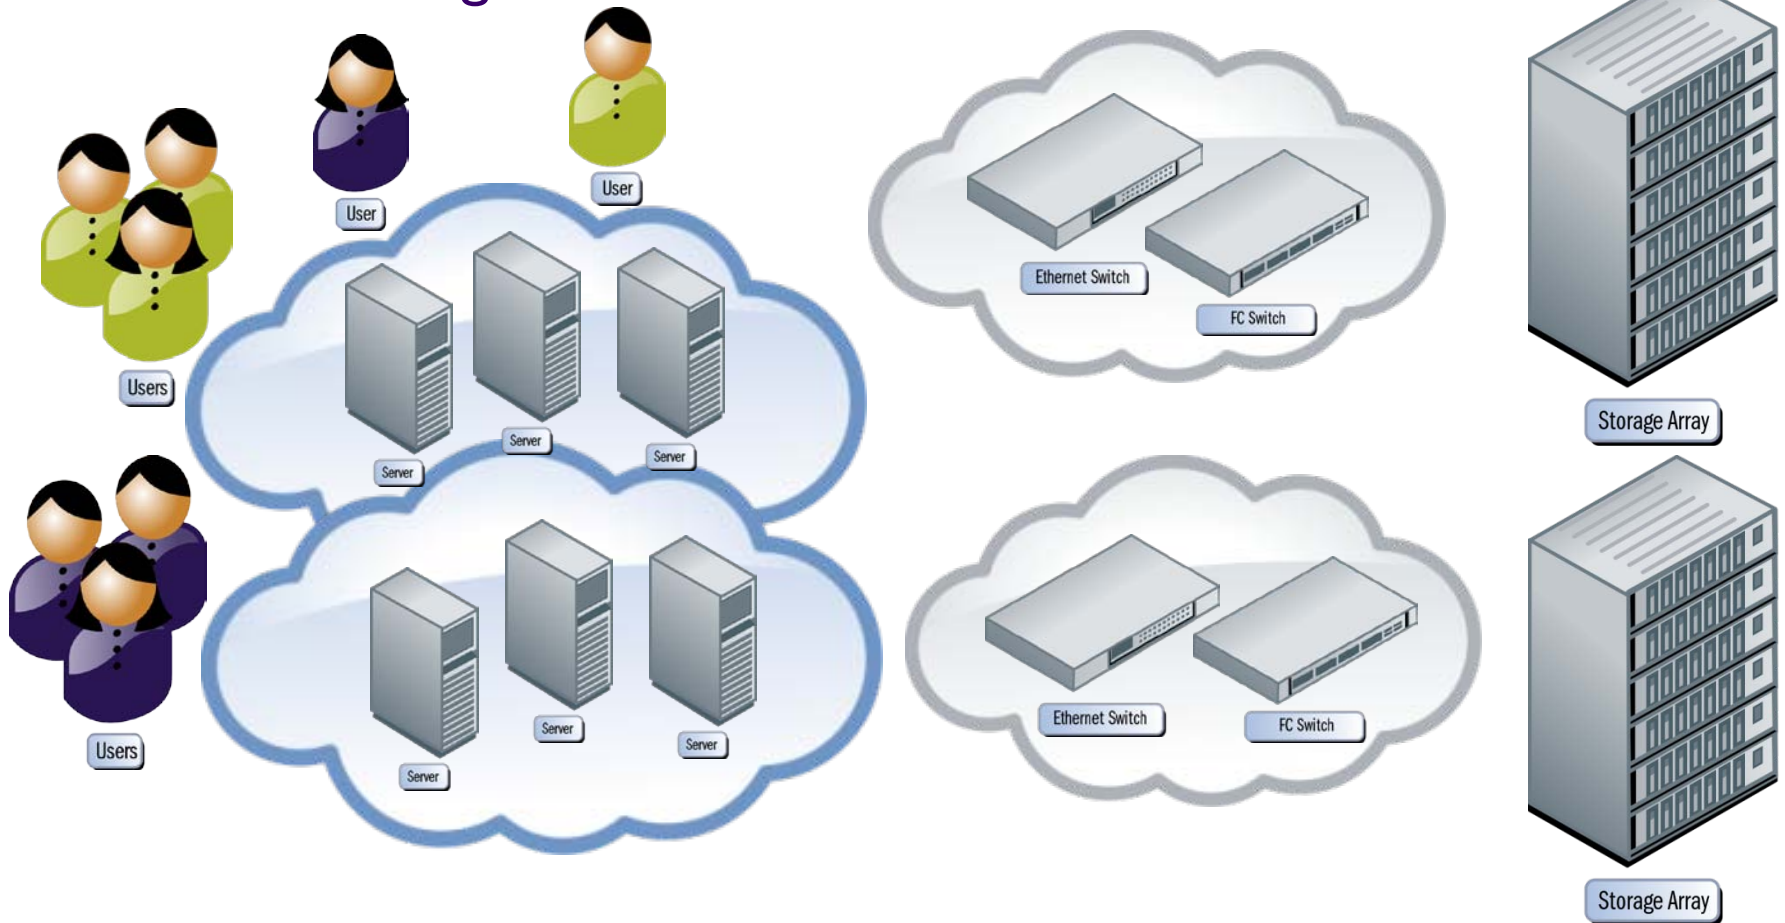
- ◆ Physical compute, network, storage management
- ◆ Virtual machine resource allocation, reservation, and limit
- ◆ Resource competition and performance

➤ Solution Scaling

- ◆ Scaling a virtual solution to multiple users
- ◆ Ease of deployment and management

Multi-user Virtual Solutions on a Shared Infrastructure

- Virtual solutions can be deployed on a shared environment consisting of shared server pools, shared storage pools, and the interconnecting fabric.



Advantages of Shared Infrastructure

- Better resource utilization.
- Centralize resource management and configuration.
- Flexible performance targets through QoS.
- Manage with familiar compute, network, and storage tools.

Disadvantages of a Shared Infrastructure

- High deployment cost and long implementation lead time.
- Need secure multi-tenant environment to address security concerns.
- Require QoS guarantee to manage resource contention.
- Potential large scale impact without high performance and high availability.




Check out SNIA Tutorial:

Understanding High Availability in the SAN

Example Physical Solution Architecture

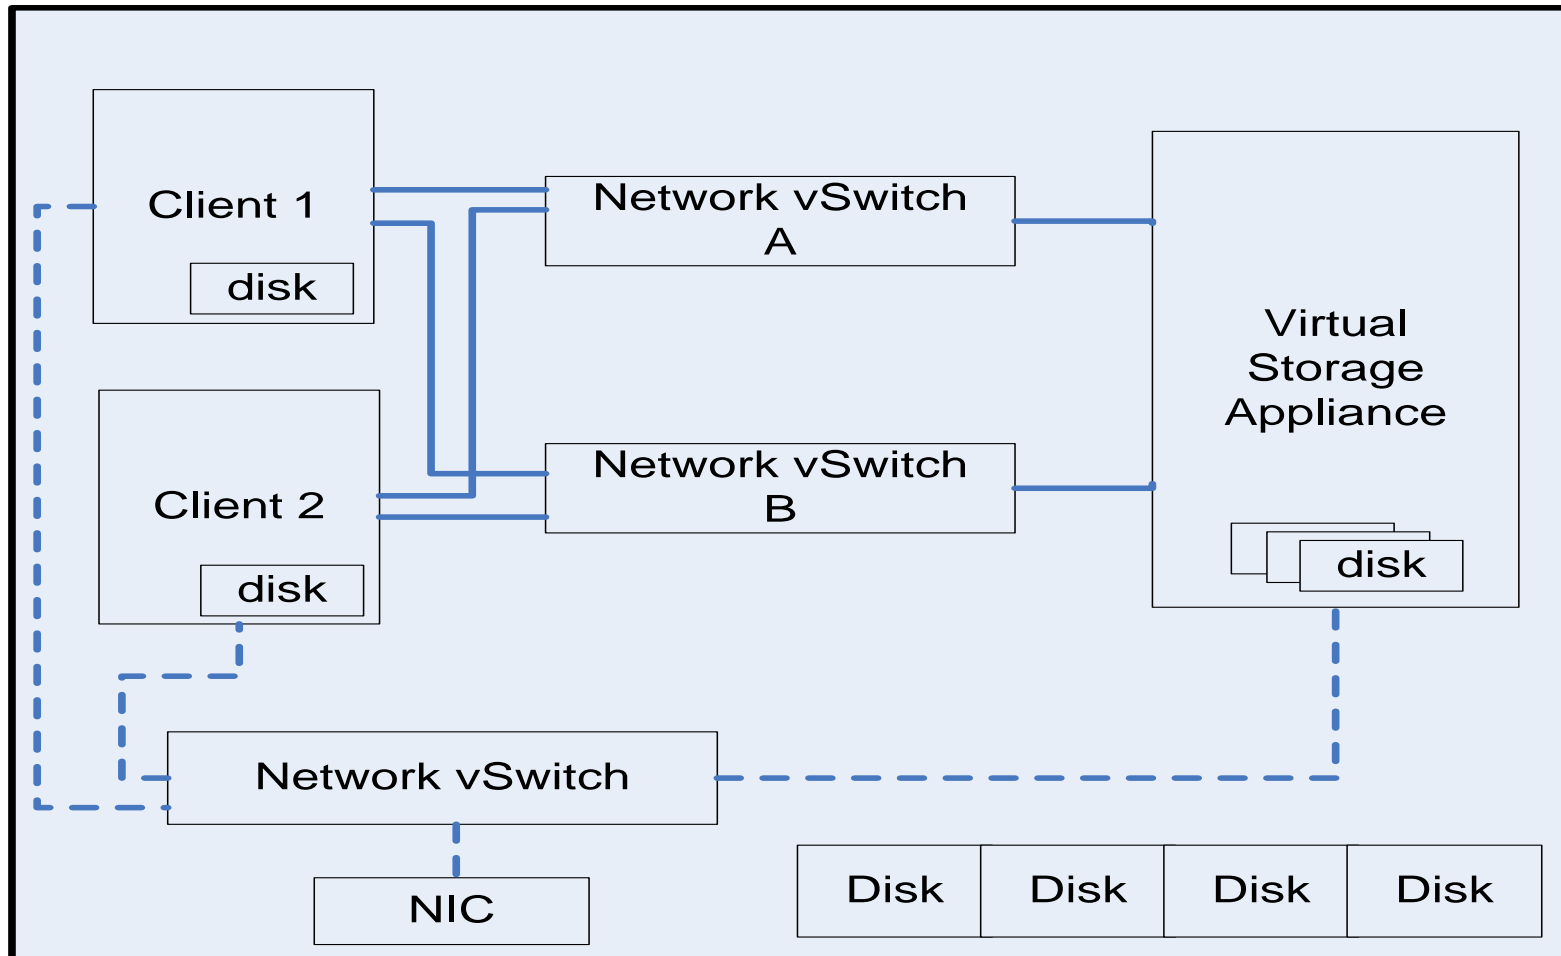
- Virtual solution architecture can be the same as the physical solution architecture.



Example Virtual Solution in a Box

Block Diagram

- Virtual clients are connected to virtual storage appliance via virtual switches, similar to how physical solution is constructed.




General Steps for Building a Complete Virtual Solution in a Box

- 1. Obtain a server with appropriate resources for the solution.
- 2. Perform RAID configuration if needed.
- 3. Install hypervisor on the server and create virtual switches.
- 4. Create virtual storage appliance and connect it to the virtual switches.
- 5. Create virtual clients and connect them to the virtual switches.
- 6. Configure the virtual storage appliance and the virtual clients to use the virtual storage using desired protocols.
- 7. Tune and test the virtual solution as needed.
- 8. Add additional virtual machines for packet capture troubleshooting or to perform network address translation if desirable.


Providing a Complete Virtual Solution in a Box

- Solution Components: virtual server, virtual switch, and virtual storage appliance.




Serving Multiple Virtual Solutions from an Isolated Box

- Solution Components: more virtual servers, more virtual switches, and more virtual storage appliances.



Boxes with Varying Capabilities

- Differences are in the available resources: CPU speed / cores, amount of memory, number of HBAs, storage type and capacity.



Performance Tuning for Virtual Solution in a Box - I

➤ Compute Resources

- ◆ Utilize hypervisor resource management features to statically or dynamically manage the resources.
- ◆ Increase the client performance by increasing the number of virtual CPUs.
- ◆ Increase the reservation of the CPU resources to obtain guarantee performance.
- ◆ Increase RAM size and set memory reservations to avoid memory swapping and ballooning.
- ◆ Set CPU resource limits on clients to protect other clients from performance fluctuations.

Performance Tuning for Virtual Solution in a Box - II

➤ Network Resources

- ◆ Increase the number of virtual network adapters to explore multipathing behavior.
- ◆ Apply network I/O control to guarantee or limit networking bandwidth.
- ◆ Utilize distributed virtual switch offerings for additional features and ease of management.
- ◆ Add PCI pass-through / SR-IOV network devices to reduce virtualization overhead.



Check out SNIA Tutorial:

PCI Express Impact on Storage Architectures and Future Data Centers

Performance Tuning for Virtual Solution in a Box - III

➤ Storage Resources

- ◆ Apply similar techniques used on tuning the client compute resources.
- ◆ Deploy storage I/O control for storage I/O bandwidth allocation.
- ◆ Utilize higher performance and redundant local storage such as SSD and RAID to achieve performance and redundancy objectives.
- ◆ Utilize shared storage to achieve reliability, availability, and performance objectives.



**Check out SNIA Tutorial:
Storage Virtualization I
What, Why Where and
How?**



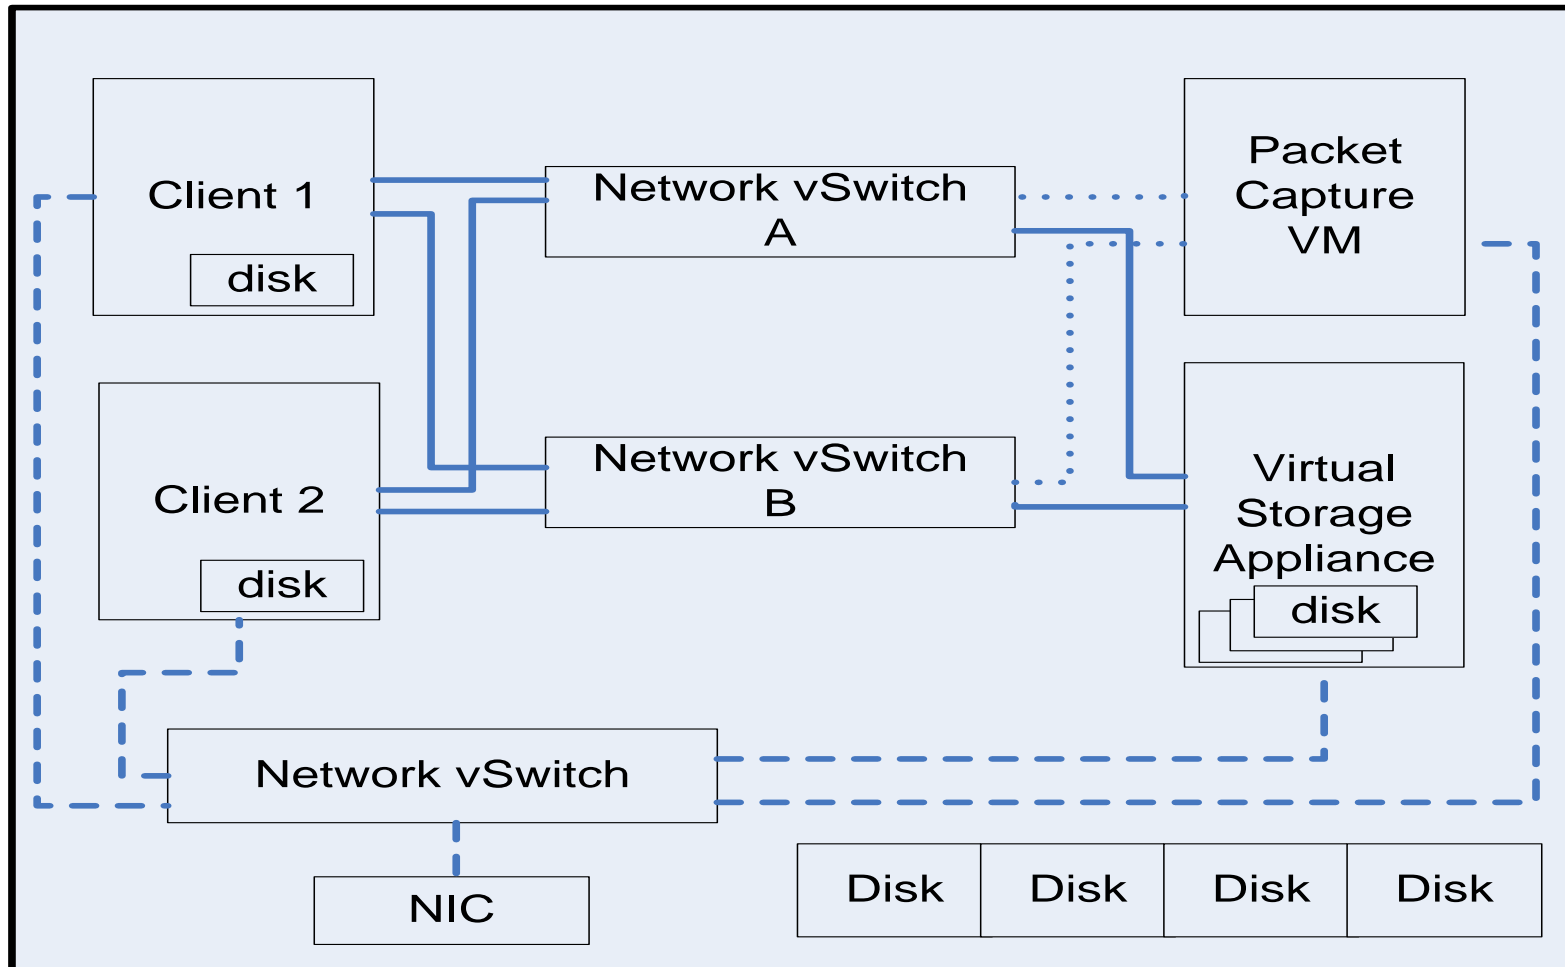
**Check out SNIA Tutorial:
Storage Virtualization II
Effective Use of
Virtualization**

Troubleshooting for Virtual Solution in a Box

- Review and analyze logs similar to a physical environment.
- Network packet tracing tools typically used in troubleshooting a physical network environment can be used: wireshark, tcpdump.
- Utilize a pre-configured packet capture virtual machine (pcap VM) helps speed up the troubleshooting process.
- Connect the pcap VM to the virtual switches for traffic monitoring.
- Provide sufficient virtual resources to the pcap VM to ensure proper traffic capture.
- Tune the pcap VM network buffers to avoid packet drop.
- Utilize RAM disk for high-bandwidth packet capture.

Troubleshooting Virtual Solution with Packet Capture VM

- Packet capture VM can be used to monitor information flowing between clients and virtual storage for troubleshooting.



- Configure interface for jumbo frame if needed.

```
#ifconfig eth1 mtu 9000
```

- Save captured trace in pcap format for trace review.

```
#tcpdump -i eth1 -w /tmp/tcpdump.pcap -c 1000000
```

```
tcpdump: WARNING: eth1: no IPv4 address assigned
```

```
tcpdump: listening on eth1, link-type EN10MB (Ethernet), capture size 96 bytes
```

```
1000000 packets captured
```

```
1623785 packets received by filter
```

```
594546 packets dropped by kernel
```


- Avoid name resolution (DNS, protocol, port) with `-nn` flag.
- Save to RAM disk for faster capture `/dev/shm/tcpdump.pcap`.
- Save specific payload capture size with `-s` flag.
- Tune network buffer for less dropped packet.

```
echo 33554432 > /proc/sys/net/core/rmem_default ;
```

```
echo 33554432 > /proc/sys/net/core/rmem_max ;
```

Tuning to Reduce Dropped Packets

- Increase `rmem_default/max` can lead to reduced number of dropped packets
- IOmeter, 16K random read, 2 worker, 32 outstanding I/O, capture 1000000 packets, 300 MB/sec.



Virtual Solution in a Box Deployment Considerations - I

➤ HW Resources

- ◆ Provide sufficient hardware resources to achieve the solution objectives.
- ◆ Resources include the needed compute and storage resources for the components and the associated virtualization overhead needed by the hypervisor.
- ◆ Test the virtual solution to ensure that it is performing at the right level.

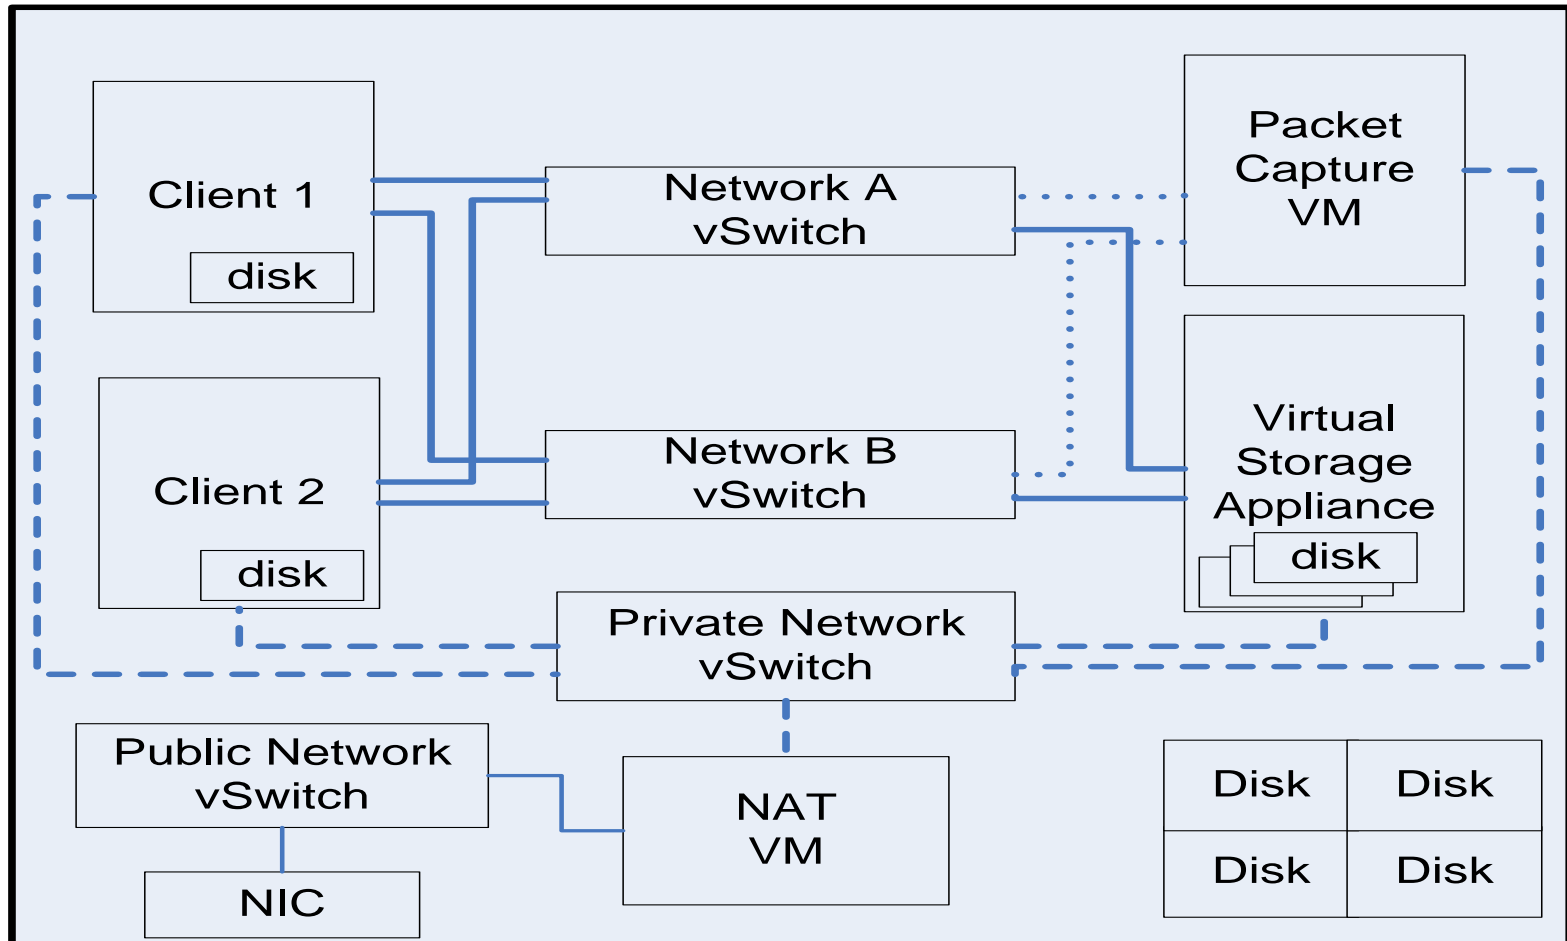
Virtual Solution in a Box Deployment Considerations - II

➤ Deployment Practices

- ◆ Deploy solutions by cloning of the environment with drive cloning if everything is in a box and on a single drive.
- ◆ Depersonalizing the solution environment allows it to be cloned and re-deployed easily.
- ◆ Utilize private IP space internally inside a solution to allow easier solution cloning.
- ◆ Utilize IP Masquerade network address translation (NAT) to allow external network access.
- ◆ Utilize shared storage LUN cloning to deploy to multiple servers.
- ◆ Utilize shared storage snapshot technology to capture and revert the environment to a known good state.
- ◆ Need special attention to provide consistent snapshot due to the possibility of caching at the various layers.


Virtual Solution with IP Masquerade Network Address Translation

- Utilize IP Masquerade network address translation for enhanced security and ease of solution cloning.



Deploying Virtual Solutions in a Box with a Hybrid Infrastructure

- Virtual solution in a box can utilize storage from a shared infrastructure.



Summary: Providing a Complete Virtual Solution in a Box

- Provide solution in a box with virtual resources.
- Manage resources to achieve desired performance.
- Apply physical environment troubleshooting techniques.
- Clone solution for additional deployments.
- Utilize shared infrastructure with hybrid deployment.
 - ◆ For large scale solution cloning and deployment.
 - ◆ To achieve reliability, availability, and performance objectives.

◆ Tcpdump

<http://www.tcpdump.org/>

◆ Wireshark

<http://www.wireshark.org/>

◆ Linux IP Masquerade HOWTO

<http://tldp.org/HOWTO/IP-Masquerade-HOWTO/>

◆ High Performance Packet Capture

<http://www.net.t-labs.tu-berlin.de/research/bpcs/>

Send any questions or comments on this presentation to SNIA: tracktutorials@snia.org

**Many thanks to the following individuals
for their contributions to this tutorial.**

- SNIA Education Committee

**Jyh-shing Chen
Joseph White**

Joshua Tseng