



Education

Advanced iSCSI Management

April, 2008

Gene Nagle, iStor Networks

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individuals may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced without modification
 - ◆ The SNIA must be acknowledged as 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.

➤ Advanced iSCSI Management

- ◆ This presentation will provide an overview of the more advanced capabilities of iSCSI storage systems and address best practices that various types of users and managers should consider in order to take maximum advantage of these iSCSI features for high availability, performance, security and ease of management.

- **Managing the iSCSI SAN for Availability & Performance**
 - ◆ Multi-pathing techniques
 - ◆ Frame sizes
 - ◆ Controller fail-over
- **Managing the iSCSI SAN for Security**
 - ◆ VLANs
 - ◆ ACL
 - ◆ CHAP
 - ◆ IP Sec and encryption
 - ◆ SSL for management interface
- **Management tools for Operation of the iSCSI SAN**
 - ◆ iSCSI targets - manage by exception
 - ◆ iSCSI initiators
 - ◆ iSNS Name Service
 - ◆ Microsoft VDS
 - ◆ SMI-S


iSCSI Availability & Performance

- Network topology and settings are vital for optimizing iSCSI availability and performance
- Make the right network configuration choices
 - ◆ Enhance availability & performance with multiple paths
 - › Choose the right multi-pathing technology
 - ◆ Choose larger frame sizes for larger I/O bandwidths
 - › Examples: Backup, video streaming
 - ◆ Take advantage of advanced features at various network levels
 - › Examples: Packet prioritization, Windows Chimney Offload



Multi-Pathing: Link Aggregation Groups

- ▶ Link-level path redundancy between network peers
 - ◆ All connections on one switch
- ▶ Increase redundancy for higher availability
- ▶ May not provide performance enhancement or (static) load balancing
- ▶ IEEE802.ad industry standard



Multi-Pathing: MCS and MPIO


- Host-centric techniques for **dynamic load balancing** and **fault tolerance** where there is more than one path between a server and a storage array.
 - ◆ **Multiple Connections per Session (MCS)**: Load balancing & path fail-over among multiple iSCSI connections within a single iSCSI session
 - ◆ **Multi-Path I/O (MPIO)**: Load balancing & path fail-over among multiple iSCSI sessions (1 connection per session) within a single I-T-L

- For fault tolerance: if a port or switch fails, the software (iSCSI driver or OS utility) can route the I/O through the other path, transparent to the application.
 - ◆ Passive fail-over/fail-back with **MCS** (session continues on remaining link(s))
 - ◆ Proactive fail-over/fail-back with **MPIO**

- For dynamic load balancing:
 - ◆ **MCS** load balancing applies to individual disks/LUNs exposed to the session
 - ◆ **MPIO** load balancing of all LUNs exposed to a shared target portal

MCS and MPIO

- ▶ iSCSI gives you multi-pathing choices
- ▶ Match the performance and availability requirements of each server



HA iSCSI RAID

- There must be ***physical*** paths from each server to each controller
- Biggest failover challenge is when using active-active controllers
- When the surviving controller takes over all I/O, it must take on the IP addresses of the target nodes presented by the failed controller (Aliasing)

Performance & Frame Size

- Frame Size is also called MTU (Maximum Transmission Unit)
- Standard Ethernet MTU = 1500 bytes
- Jumbo Frames can be anything larger than 1500
 - ◆ Normally up to 9000 bytes
 - ◆ Jumbo frames are a good complement to Gigabit Ethernet
- Best Practice Consideration
 - ◆ To use jumbo frames every device in the network must be able to handle them & must be configured properly
 - › Large MTU sizes can reduce overhead, improve throughput
 - › If I/O sizes are small, jumbo frames are under utilized
 - ◆ To mix MTU sizes, use VLANs



Managing iSCSI SAN Security

- A number of well proven techniques are available to make iSCSI SANs secure
- Most basic measure is **isolation** of the SAN from other networks
 - ◆ This is all that protects most FC SANs
 - ◆ Exception: connectivity to the LAN/WAN for management, usually via separate ports from data ports



- “Zoning” for iSCSI: VLANs
 - ◆ A virtual LAN is a technique to segregate network users on the network
 - ◆ The VLAN restricts access of specific users to a given target storage system.
 - ◆ VLAN zoning is also a way to manage traffic on the network for optimal efficiency and accelerated performance.

- VPN across an unsecured network
 - ◆ Well-established solutions for creating secure, virtual point-to-point IP bridges across un-trusted mediums such as the Internet.

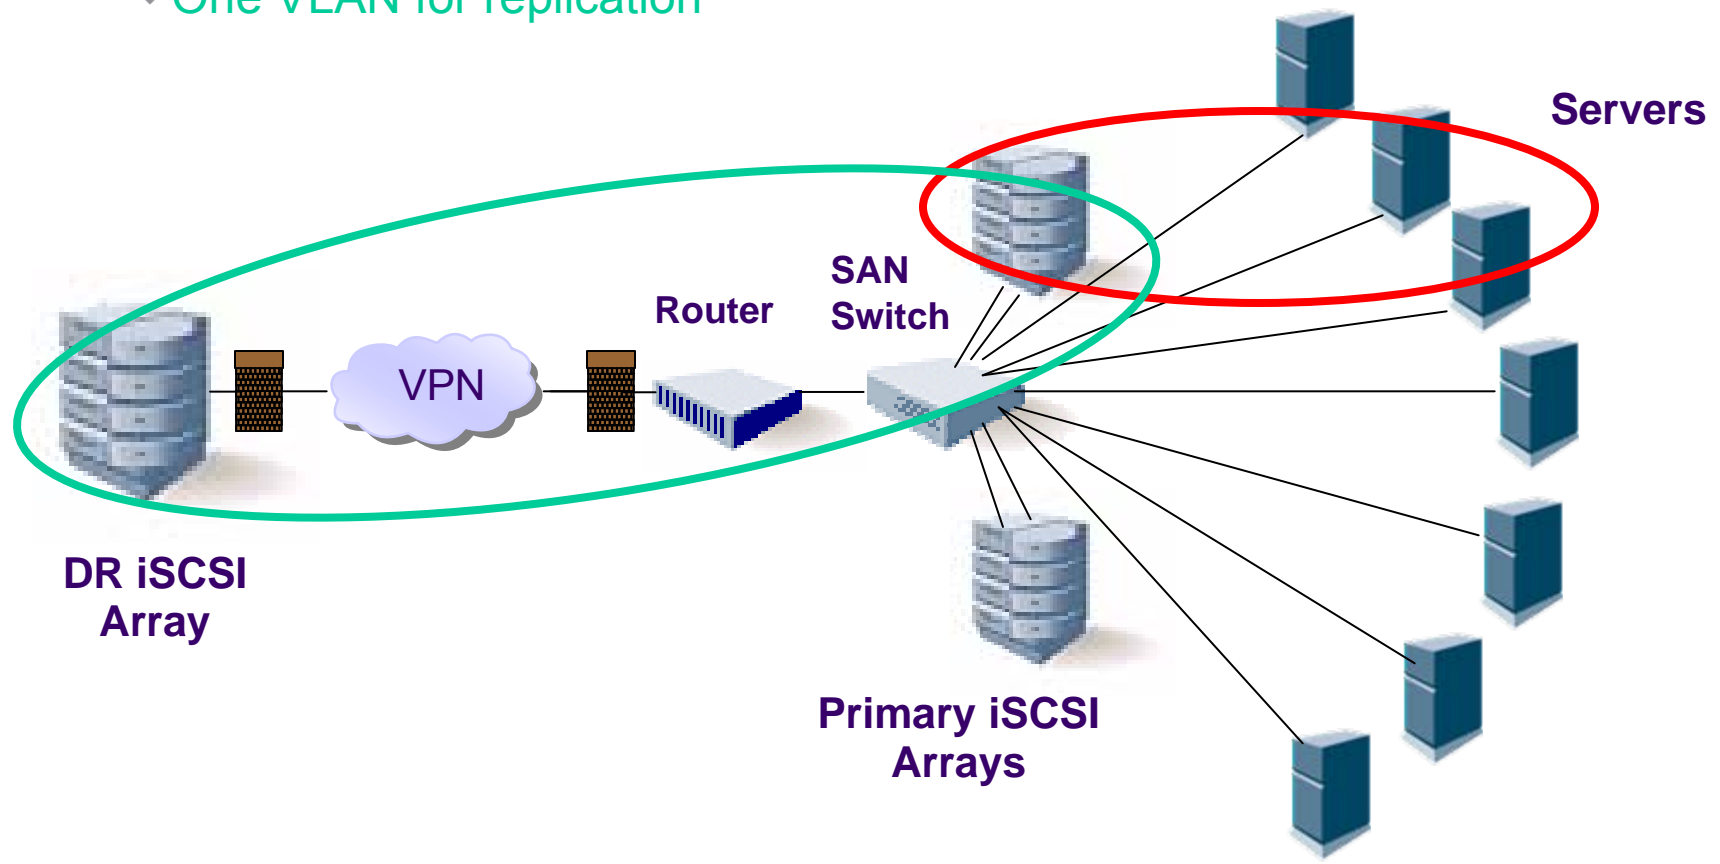
- Access Control Lists (ACLs) by the iSCSI target

- Authentication and encryption: choices....

- Secure management connections

Using VLANs and VPN for the DR Site

- ▶ LUNs being replicated can be in 2 VLANs
 - ◆ One VLAN for primary site access
 - ◆ One VLAN for replication



- ◆ **A Form of LUN Mapping & Masking:** ACL support to the volume level
 - ◆ ACLs (alone) have been proven to be insecure in many situations

- ◆ **LUN masking:** Each LUN may have access to it restricted to a specified iSCSI initiator or group of initiators

- ◆ This limits which iSCSI initiators can 'see' an iSCSI LUN. Typically, just one iSCSI initiator has access to an iSCSI LUN, preventing write collisions but also providing privacy for that LUN

- ◆ ACL security isn't sufficient when un-trusted users have root access on a system capable of accessing the target

- **Read-only volumes** can permit multiple users to all have concurrent access to the data and maintain data integrity without managing user privileges
- Check availability of this feature with your hardware manufacturer

Authentication Alternatives

- **CHAP** authentication is most common for iSCSI

- **RADIUS** is primarily used for embedded network devices such as routers, modem servers, switches

- **TACACS+ & LDAP**
 - ◆ TACACS+ is a Cisco protocol
 - ◆ Lightweight Directory Access Protocol, or LDAP is an application protocol for querying and modifying directory services running over TCP/IP
 - › LDAP natively provides no protection against sniffing or active attackers

- **IPSec & Kerberos** combine authentication and encryption

CHAP Authentication

❖ Challenge Handshake Authentication Protocol

- ◆ Based on peers (initiators, targets) sharing a secret
- ◆ One-way CHAP: only target authenticates initiator(s)
- ◆ Mutual CHAP: target and initiator authenticate each other

❖ CHAP Authentication – done during initial iSCSI login to authenticate end nodes, restrict access to targets

- ◆ Vulnerable to sniffing
- ◆ Vulnerable to message reflection attacks across multiple connections when two hashes match the password can be compromised
- ◆ Less secure than IP Sec, but less impact on iSCSI performance

❖ Best Practices for open iSCSI SAN



- ◆ At a minimum, the use of one-way CHAP authentication between iSCSI initiators and targets is typically recommended
- ◆ Consider iSNS IPsec for added security

Using IPSec

- ◆ **IP Security (IPSec)** can be used as an alternative to **CHAP** to ensure that iSCSI end points (initiators and targets) are authentic and to maintain privacy and integrity of transferred data (cryptographic integrity)

- ◆ **Authentication and Encryption with IPSec**
 - ◆ IPSec may be used to encrypt authentication and data packets on the network. A common key is set on all IP portals, allowing all peers to authenticate each other and negotiate packet encryption.
 - ◆ IPSec provides two levels of security:
 - ◆ Authenticates both the initiator and target nodes, preventing the “**man-in-the-middle**” type of attacks.
 - ◆ Encrypts the data being transferred on the network so that any network snooping that is taking place would only capture garbled data.

- ◆ **Best Practice Consideration**
 - ◆ IPSec creates significant performance overhead



Securing the Management Connection

- ◆ The iSCSI storage management interface should be secured to prevent unauthorized access from hackers.
 - ◆ Particularly true with web-based configuration tools that can be accessed from anywhere.

◆ Best Practices:



- ◆ Use Secure Socket Layer (SSL) Encryption for Management Interface Security
- ◆ Use VPN, even within the LAN



- Audit iSCSI devices and networks to assess risk
- Security needs to be addressed end to end (host – network – device)
- Configure a network topology that minimizes risk of unauthorized access to or modification of data as it traverses the network;
 - ◆ Isolate the SAN from other networks
 - ◆ Avoid DHCP on the SAN
 - ◆ Virtual LANs (VLANs)
- Extra precautions for distributed data center and remote replication (DR) applications
 - ◆ WANs require stronger security measures
- The level of security that you can set for a storage subsystem depends on the hardware manufacturer.
 - ◆ Not all subsystems support all levels of iSCSI security
 - ◆ Contact your hardware manufacturer

- Take advantage of advanced features in **iSCSI targets**
 - ◆ High level management by exception

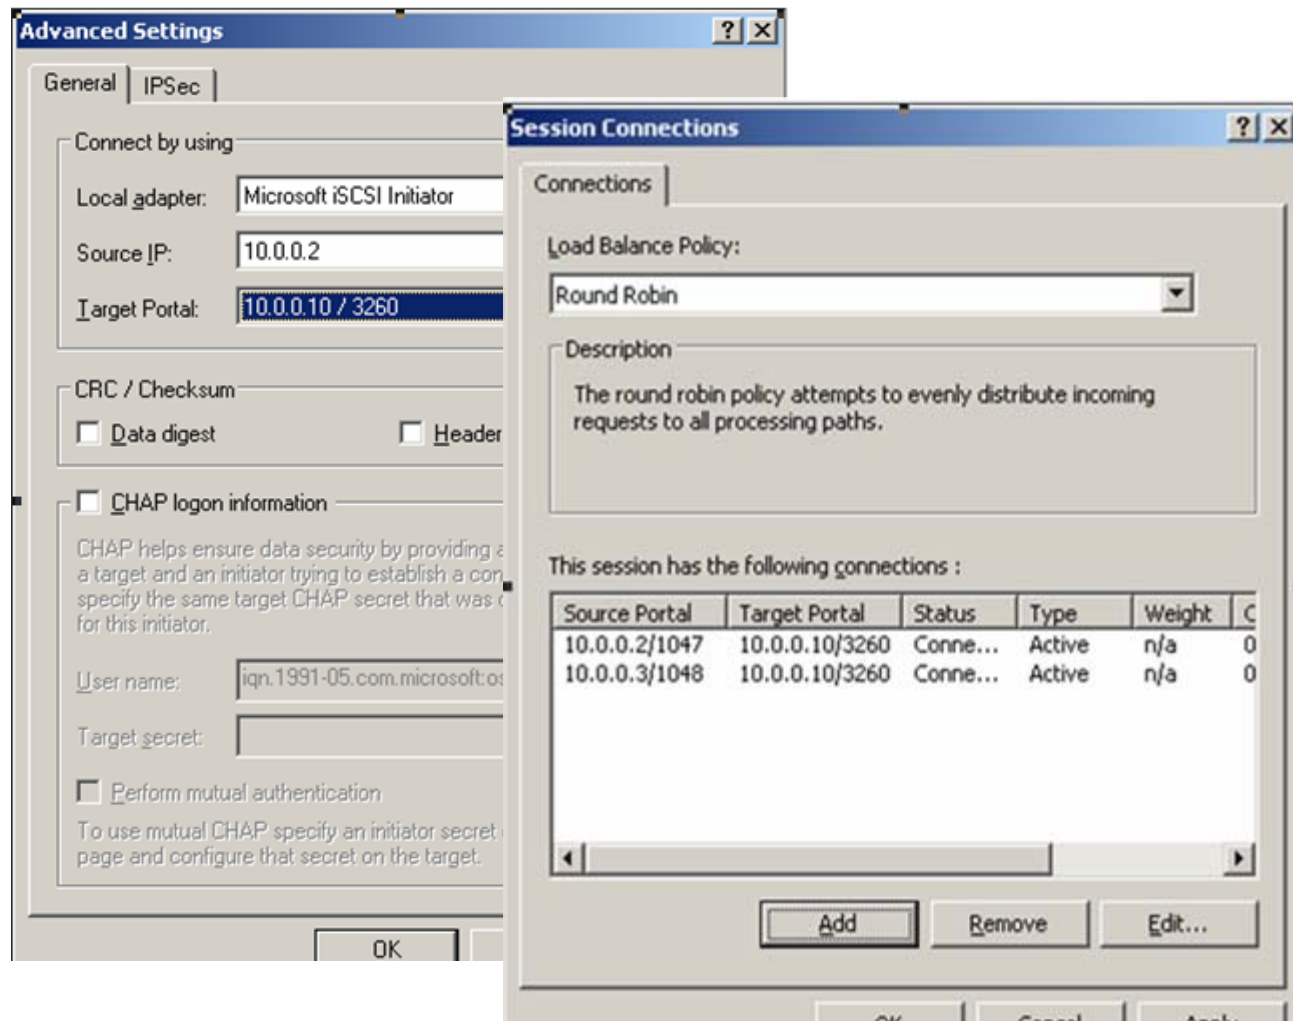
- Take advantage of advanced features in **iSCSI initiators**
 - ◆ Microsoft iSCSI initiator
 - ◆ Hardware initiators

- Consider using **iSNS** name service

- Take advantage of advanced management tools
 - ◆ Microsoft Virtual Disk Service (VDS)
 - ◆ SMI-S

Sample iSCSI Initiator

- Use Advanced Settings for security and performance enhancement



Using iSNS iSCSI Name Service

➤ iSNS and Discovery Domains

➤ iSCSI Name Services

- ◆ Groups of iSCSI initiators and targets in Domain for logical segmentation
- ◆ iSCSI initiators and targets register with the iSNS Server (similar to DNS)
- ◆ Initiators only discover the targets they are eligible to use



➤ iSNS Best Practice

- ◆ Use unique names not default naming for Domains
- ◆ Important at iSNS registration to move nodes out of the Default Domain Pool for enumeration and access control
- ◆ Control specific port for iSNS Server to prevent (Man-In-The-Middle) attacks from a fake iSNS Server

Storage Management with iSCSI

➤ SMI – Storage Management Initiative

◆ Interoperable Management Interface

- > Based on Common Information Model (CIM) from Distributed Management Task Force (DMTF)

- > Runs over HTTP

◆ Profiles implement IETF iSCSI model

- > Initiator

- > Target

➤ IMA – iSCSI Management API

◆ Host Driver/HBA Interoperability

◆ ANSI 411

SMI Discovery

- Target Nodes
 - ◆ SCSI targets
 - ◆ World Wide Names
- Network Portals (IP Addresses used for iSCSI)
- Access Paths
 - ◆ Logical Unit Numbers
 - ◆ Read/Write Permissions
- Sessions and Connections
 - ◆ Current Connections
 - ◆ Max Connections
 - ◆ Initiator Name
 - ◆ Negotiated Parameters
- Statistics
 - ◆ Protocol Data Units
 - ◆ Logins
 - ◆ Errors

SMI iSCSI Configuration

➤ Target Creation

- ◆ iSCSI Nodes
- ◆ iSCSI Network Portals (IP Addresses)

➤ Access Path Creation

- ◆ Volumes to Hosts
- ◆ Logical Units and Permissions

- Please send any questions or comments on this presentation to SNIA: trackstorage@snia.org

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

- SNIA Education Committee

**Gene Nagle
Jay Kramer
Scott Baker**