



THE ART OF SFTI

A FINANCIAL INDUSTRY NETWORK

MPLSCON 2006
NEW YORK, NY
MAY 23, 2006

NAISHEN WANG
SENIOR MANAGER
COMMUNICATION PLANNING
SECURITIES INDUSTRY
AUTOMATION CORPORATION

Disclaimer

COPYRIGHT NOTICE

COPYRIGHT © 2006 BY THE SECURITIES INDUSTRY AUTOMATION CORPORATION (SIAC). ALL RIGHTS RESERVED. EXCEPT AS PERMITTED UNDER THE UNITED STATES COPYRIGHT ACT OF 1976, NO PART OF THIS DOCUMENT MAY BE REPRODUCED OR DISTRIBUTED IN ANY FORM OR BY ANY MEANS, OR STORED IN A DATA BASE OR RETRIEVAL SYSTEM, WITHOUT THE PRIOR WRITTEN PERMISSION OF SIAC.

PROPRIETARY NOTICE

DISCLOSURE OF THE CONTENTS OF THIS DOCUMENT IN WHOLE OR IN PART TO ANY OTHER PARTIES WITHOUT THE PRIOR WRITTEN CONSENT OF SIAC IS EXPRESSLY PROHIBITED.

BRAND NAMES AND/OR TRADEMARKS


BRAND NAMES OR PRODUCTS CITED IN THIS DOCUMENT MAY BE TRADE NAMES OR TRADEMARKS. WHERE THERE MAY BE PROPRIETARY CLAIMS TO SUCH TRADEMARKS OR TRADE NAMES, THE NAME HAS BEEN USED WITH A INITIAL CAPITAL. REGARDLESS OF THE CAPITALIZATION USED, ALL SUCH USE HAS BEEN IN A EDITORIAL FASHION WITHOUT ANY INTENT TO CONVEY ENDORSEMENT WHATSOEVER OF THE PRODUCT OR TRADEMARK CLAIMANT. SIAC EXPRESSES NO JUDGMENT AS TO THE VALIDITY OR LEGAL STATUS OF ANY SUCH PROPRIETARY CLAIMS.

ENGINEERING SERVICES DISCLAIMER

INFORMATION CONTAINED IN THIS DOCUMENT IS BELIEVED TO BE ACCURATE. HOWEVER SIAC DOES NOT GUARANTEE THE COMPLETENESS OR ACCURACY OF ANY OF THE PUBLISHED INFORMATION. THIS WORK IS PUBLISHED WITH THE UNDERSTANDING THAT SIAC IS SUPPLYING INFORMATION, BUT NOT ATTEMPTING TO RENDER ENGINEERING OR OTHER PROFESSIONAL SERVICES. IF SUCH SERVICES ARE REQUIRED THE ASSISTANCE OF THE APPROPRIATE PROFESSIONAL SHOULD BE SOUGHT.



Agenda



- Introduction
- Design Objective
- Design Overview
- Multicast Design



What is SFTI?


■ Secure Financial Transaction Infrastructure

- Pronounced “*safety*”
- A direct result of the September 11, 2001 impact on leased line connectivity throughout the USA
- Developed to help the **financial services industry** substantially improve the **resilience** of **telecommunications connectivity**
- Spans multiple cities (New York, New Jersey, Chicago, Boston, and Philadelphia)
- Over 700 member firm connections
- Web site: <http://sfti.siac.com>




SFTI Design Objective

- No single point of failure
 - Fiber, access point, data center, router, link
- Resiliency and fast recovery
- Low end to end latency
- Redundant Multicast delivery





SFTI High Level Diagram





SFTI Connectivity Requirement

- Each customer / service provider required to connect to a minimum of two Access Centers
 - At least one connection to each logical network
- Only Ethernet connectivity is accepted
 - 100BT or Gigabit
 - Customer can co-locate its access routers in SFTI Access Center
 - SFTI partners with extranet to handle other type of connections
- Customer edge router is expected
 - EBGp for unicast
 - PIM for multicast

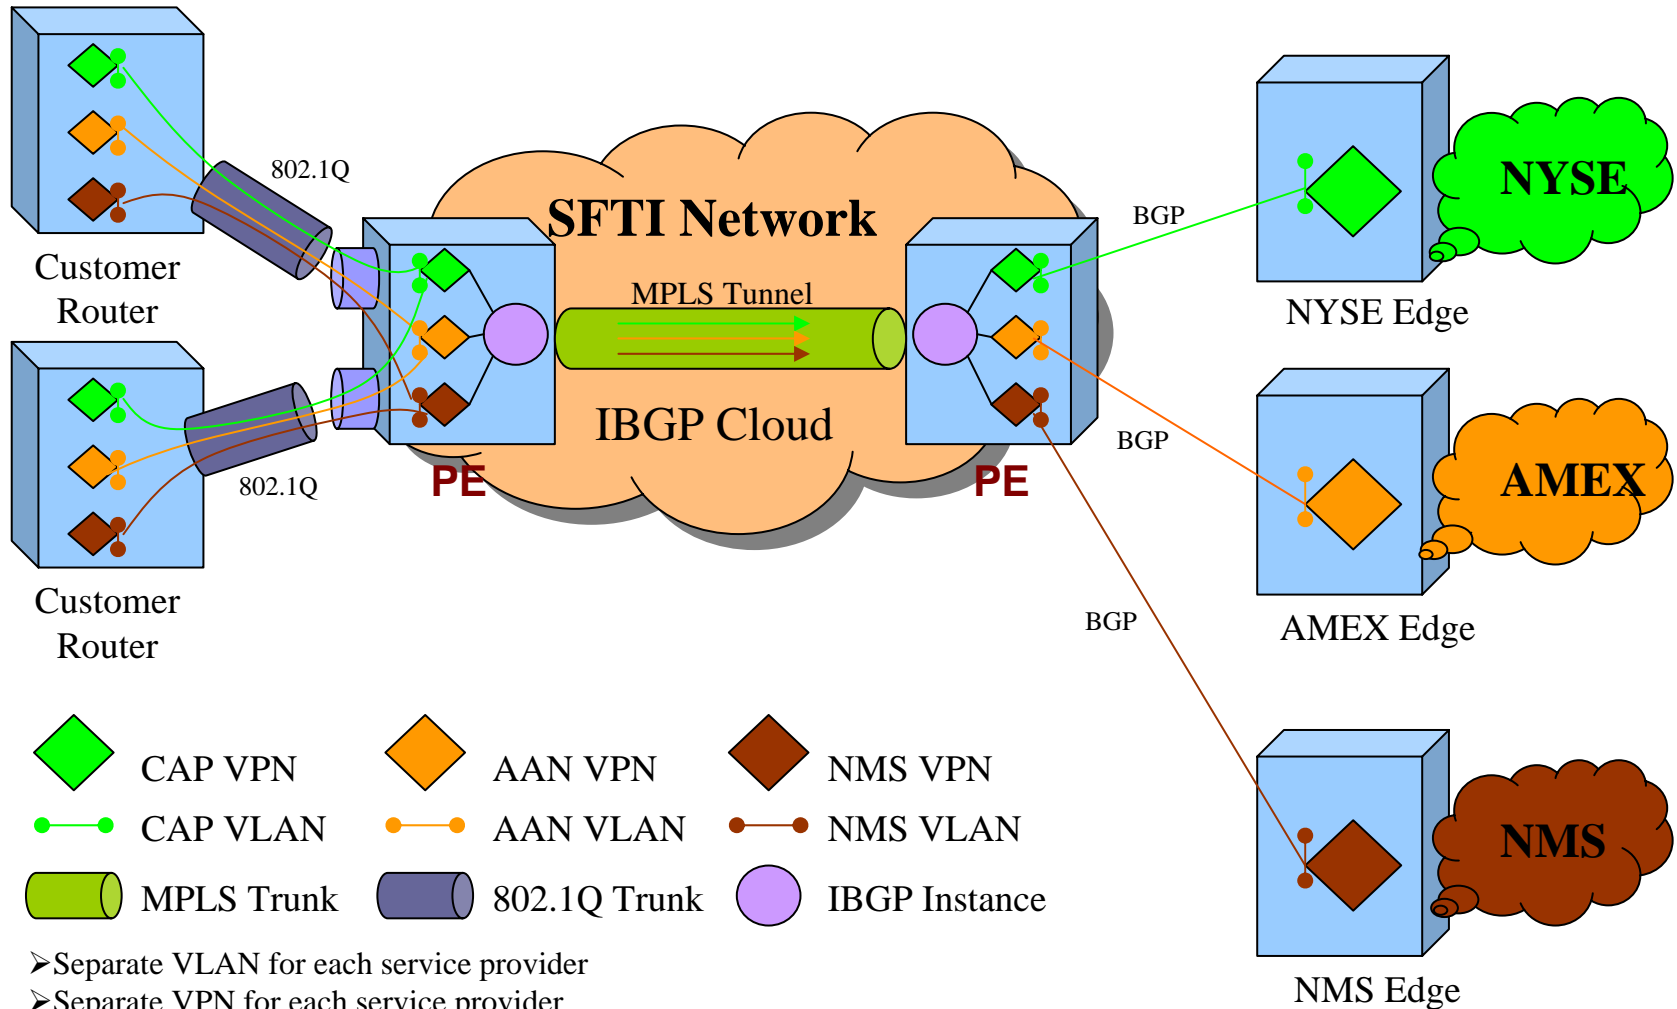


SFTI Design

- Utilizes MPLS Layer 3 VPN technology (*RFC 2547bis*) to create virtual backbones for financial service providers
- Uses VLAN trunking (IEEE 802.1q) to connect to the customers
- One VLAN associates with one service provider's virtual backbone



SFTI Network Overview





Multicast in Financial Industry

- The benefit of using Multicast as a financial market data distribution method:
 - Efficient bandwidth usage
 - Minimum server resources
- Industry-level multicast applications on SFTI
 - Consolidated Trade and Quote (CT & CQ)
 - Options Price Reporting Authority (OPRA)
 - NYSE Openbook



Multicast Requirement

■ Highly reliable

- Every **single** packet counts
- Very different from delivering video

■ Highly available

- Quick network recovery is not sufficient
- If a failure occurs in transit network, Multicast packets are not recoverable

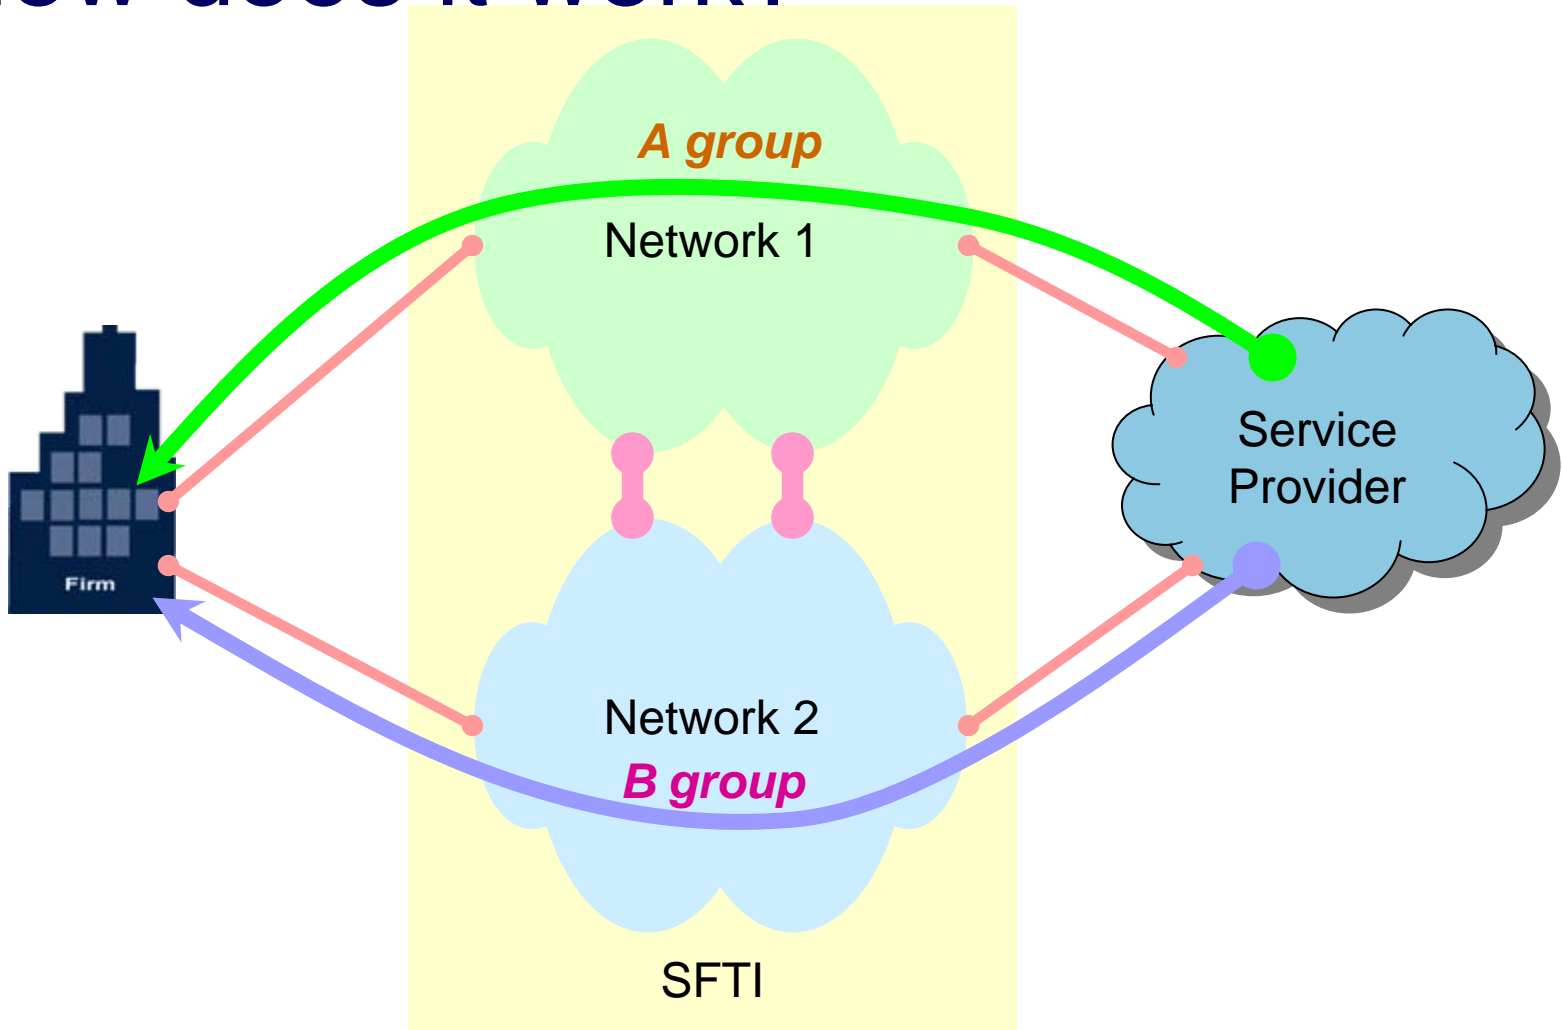


Multicast Redundancy

- SIAC is the original inventor and patent holder for providing Multicast redundancy with dual multicast streams
 - Two multicast streams with the same data, delivered to a single receiver simultaneously via two diverse network paths
 - Upon receiving, receivers compare both streams and use the best available
 - *US Patent No. 6,408,000, Multicast Data Distribution System*



How does it work?





Original Multicast Design

- **SFTI delivers NMS, NYSE and AMEX services to the financial industry**
 - **Uses global routing table**
 - **All Multicast source networks are running DVMRP**
 - **Multicast over Layer 3 VPN technology is still in its infancy**



SFTI B2B services

- Delivers unicast and multicast services for other financial market centers / content service providers to SFTI customers



Limitation of Original Model

- Operational changes for one Multicast service can affect other Multicast service providers
 - RP and multicast group association
- IP address and multicast group assignment



Multicast Design for SFTI B2B

- Multicast over Layer 3 VPN via GRE encapsulation
 - Each Market Center/Content Service Provider will have its own virtual backbone on SFTI
 - Each Market Center/Content Service Provider runs its multicast applications in the same virtual backbone as its unicast traffic



How does it work?

- Service provider runs PIM sparse mode with SFTI PE
 - Service provider requires its own RP for both SFTI and customers
- SFTI runs PIM to all customers
 - SFTI provisions customer to use Multicast services with PIM sparse mode
- Multicast VPN via GRE Encapsulation (*draft-rosen-vpn-mcast-06.txt*) on SFTI backbone

**GRE = Generic Routing Encapsulation*



MVPN via GRE Encapsulation

- Multicast GRE tunnels
 - PIM messages among SFTI PE routers
 - Actual Multicast packets
 - Via global routing table
- Utilize the L3VPN virtual backbone for each service provider
 - Source routes are distributed in service provider's virtual backbone
- Dedicated PIM domain for each service provider
 - Individual RP configuration
 - Individual PIM neighbor

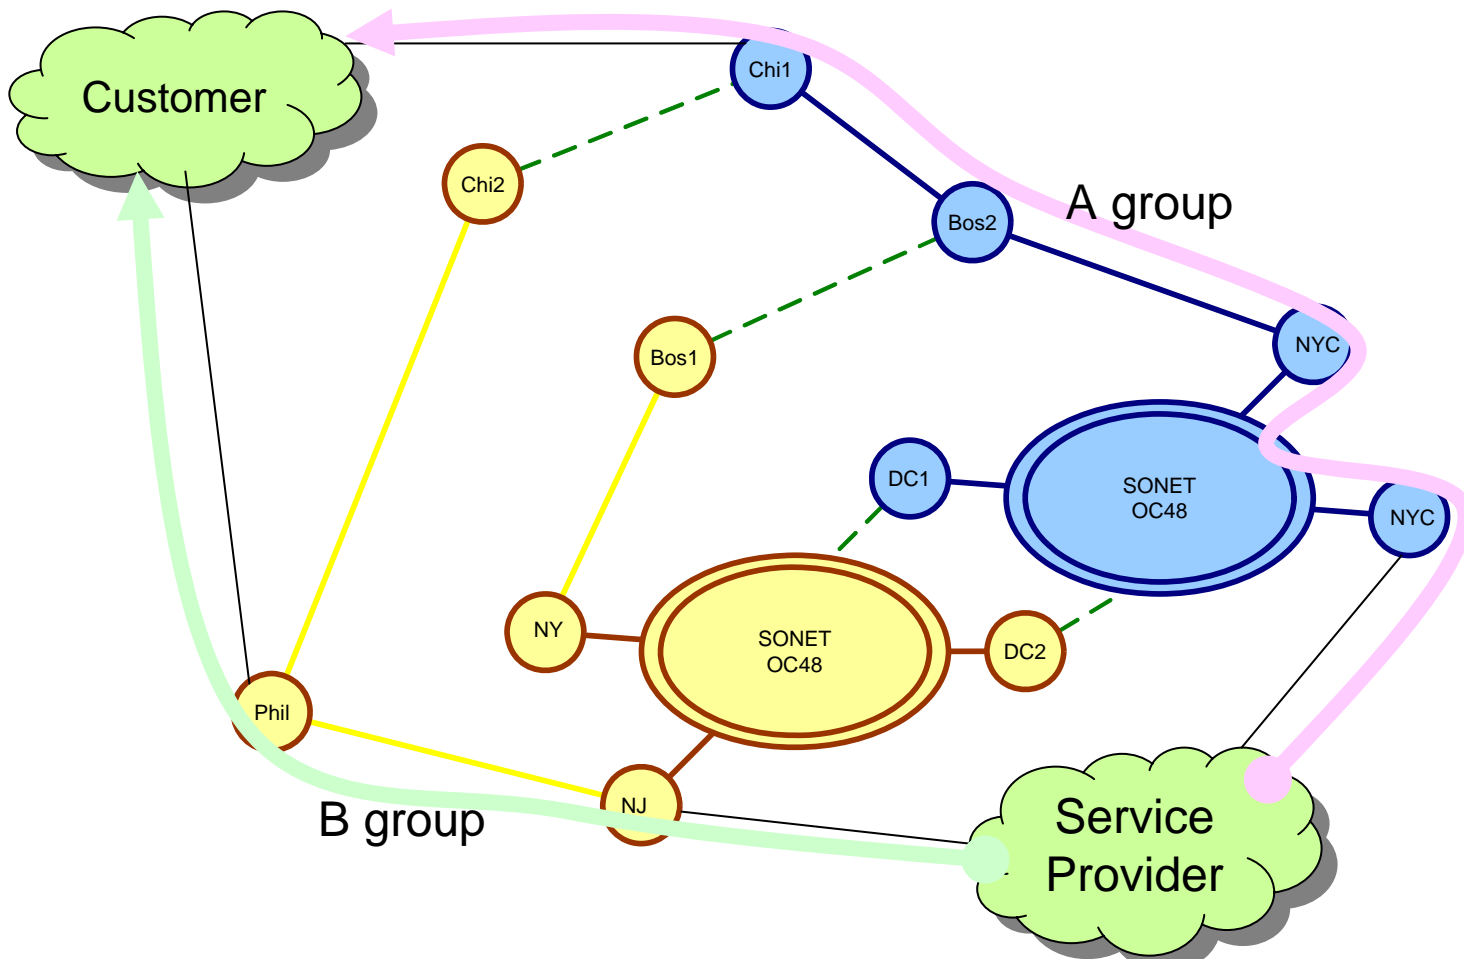


Multicast Redundancy

- Routing metric used to steer traffic flows within the same logical network as preferred
 - Raise the OSPF metric for the links between the two logical networks
- Any single point of failure in the network affects only one of the two multicast groups
 - Data gets through uninterrupted



Multicast Flow





Future Direction

- **Current limitation**
 - GRE encapsulation and de-capsulation introduce additional latency and overhead
 - Lack of PIM RP support
 - MVPN MSDP support on PE routers
 - Troubleshooting complicity
 - Lack of PIM status on virtual tunnel interface



MVPN via P2MP LSP?

- Is MVPN via P2MP LSP the right solution?
 - Vendor support?
 - How long?
 - Industry direction?





Thank You!

- Questions?
- Contact info:
 - Naishen Wang
 - nwang@siac.com

