

ntt.net



## *A Day in the Life of a Security Professional*

Peter Schoenmaker

[pds@ntt.net](mailto:pds@ntt.net)

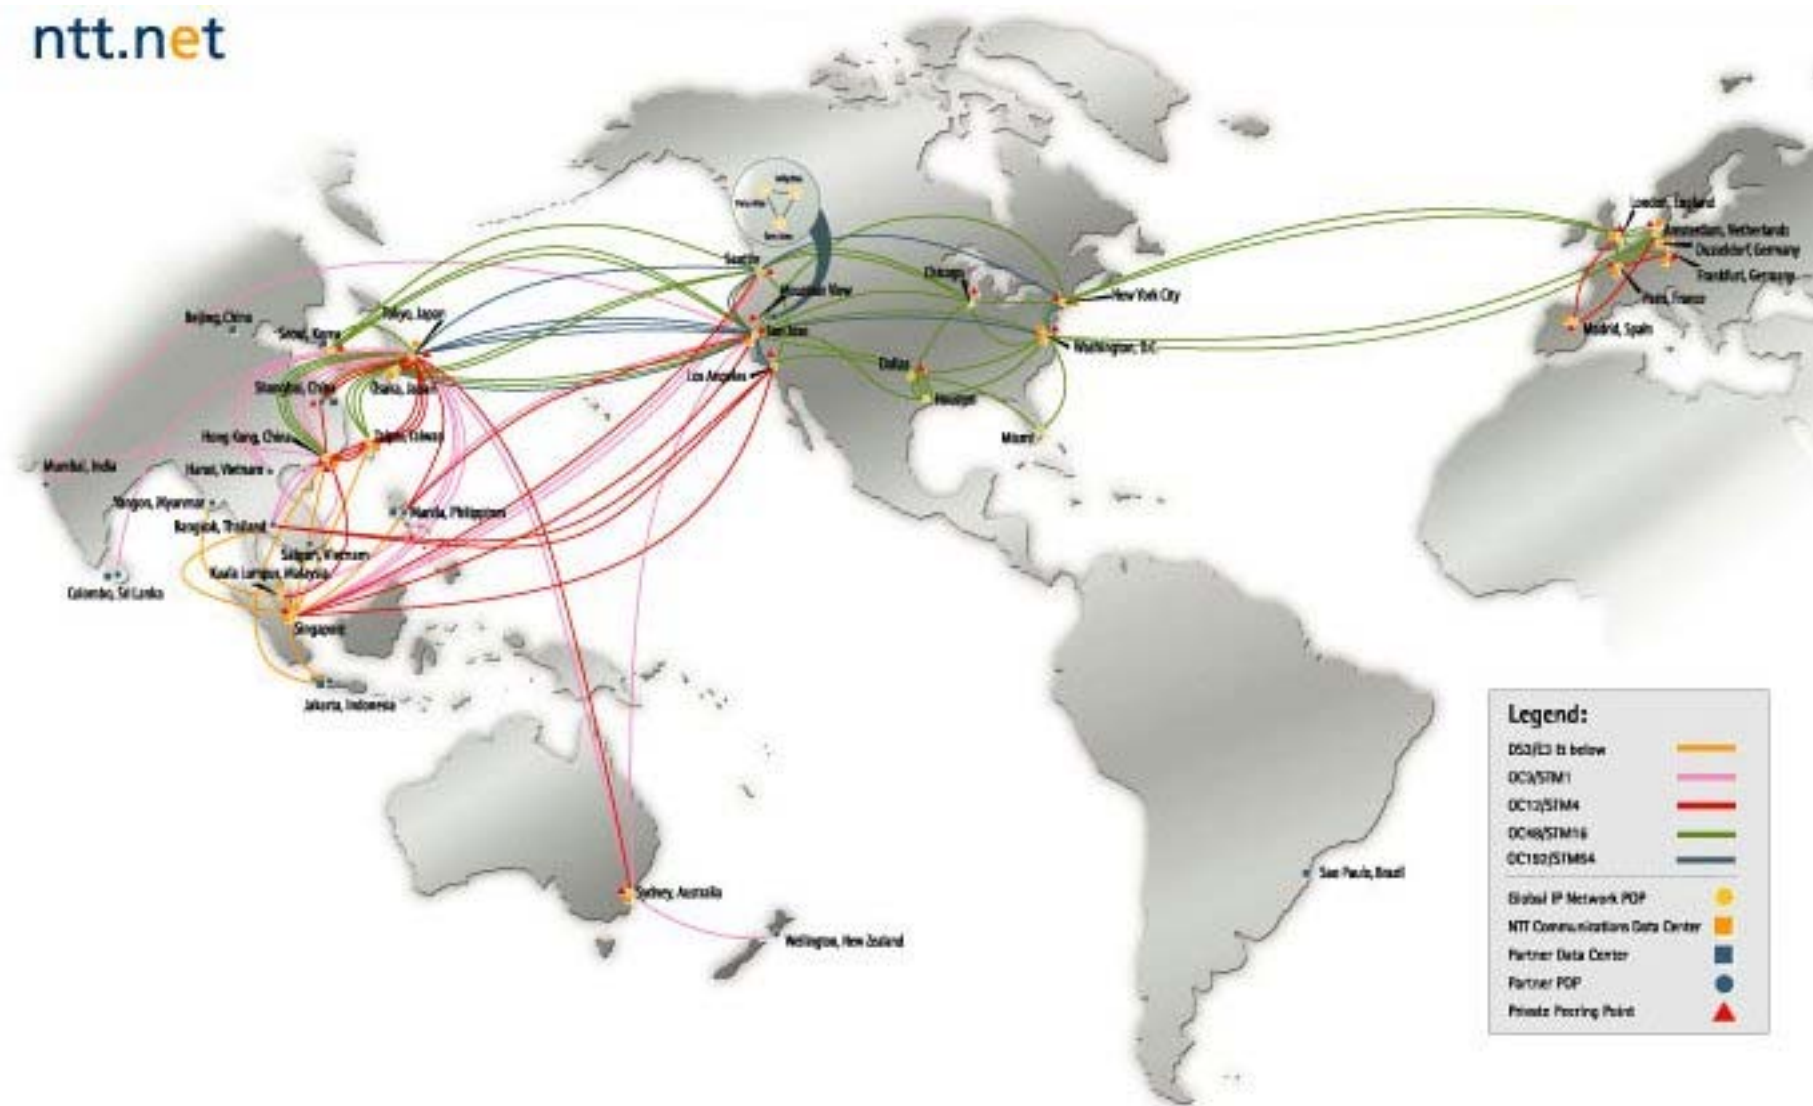
March 2nd, 2006

NTT Communications Corporation

- Peter Schoenmaker
  - Work for NTT Communication's Global IP Network (AS2914)
  - Member of the IP Network Development Group (USA)
    - Working on Network design and architecture
- NTT Communication's Global IP Network
  - Spans 4 continents, Asia, Australia, Europe, and North America
  - STM16, and STM64 network
  - Target large customers E3/DS3 or above
    - Other divisions handle smaller connections
- Disclaimer
  - The topics in the presentation are a collection of ideas from various colleagues

# NTT Communications Global IP Network

ntt.net



## *Security in the Real World*

- Security has to be planned from the beginning
- K.I.S.S. Keep it Simple and Straightforward
- Security vs. Usability
  - The only secure router is one not connected to the network
- Cost of Security
  - The price of services is dropping
  - The cost of service is dropping at a rate less than the price
  - Security can be expensive
- Users and Ability
  - Some users cannot register routes in an IRR
  - Some users do not know what a router is

- Control Plane Security
- Management Plane security
- Data Plane security
- Hardware
- Operations

## *What is the Control Plane?*

- Protocols used to program the forwarding hardware
  - BGP, OSPF, ISIS, RIP
  - LDP, RSVP
  - PIM, MSDP, IGMP
- Configures the router to forward packets
  - Microprocessor
  - ASIC

- Protocol Signaling security
  - ACL filtering
  - MD5 authentication
  - Encryption
- Route announcements
  - Route filtering
    - Stable routing
  - Authoritative announcement
  - Verified route announcement
- Routing table
- Isolate the control plane

- Filter everything on the router
  - More difficult to implement without automation
  - Filter
    - src address
    - port
    - protocol
    - `ip access-list 101 permit tcp host 192.168.1.4 any eq 179 ! BGP`
- Some vendors provide commands to more easily filter the Control Plane/Management Plane
  - Juniper: apply a filter to the lo0 interface
  - Cisco: Control Plane Policing

- eBGP md5 authentication used only recently
  - Password never changes
  - Operationally difficult to change
    - BGP Implementation updates coming
- BGP over IPSec

- NTT Communications Global IP Network, requires all customers to register routes in an IRR
  - Customers can register any route originated from any AS, or add any ASN to their as-set
    - BGP prefix-list filter is automatically updated
  - Customer route announcements based on prefix-list
    - No AS Path filtering
      - Origin AS data is lost during generation of the prefix-list filter
        - Origin AS may not be accurate
      - Customers have leaked routes
  - NTT Communications Global IP Network accepts any route announced by the customer (which is registered in the IRR)
    - Except Bogons
  - Routes, and AS-Sets are rarely cleaned up
  - Protects against mistakes
    - Malicious users could register any route, as many routes, as they want
    - Stale objects allow for mistakes (PANIX)

## *IPv6 BGP Prefix filtering*

- Large problem today
- Little or no filtering done
- Many carriers still provide free transit on public peering sessions
- Some implementations take up more FIB space

- Filtering from the IRR does not stop route hijacking
  - Difficult to determine the difference between a legitimate route and a hijacked route
- Solved with more specific routes
  - If 192.168.0.0/20 is hijacked, the owner will announce 192.168.0.0/21, and 192.168.8.0/21
- Relatively low impact on the user experience

- Some routes never change
  - AS2914 has announced 129.250.0.0/16 for 12 years
  - Should stable routes have a heavier weight vs. a new announcement?
- Some connections should never announce certain ASNs
  - Joe's T1 connection should never announce a route with level3's ASN in the path
- Problems
  - Mergers and sales
    - Verizon and MCI
    - Networks get combined, ASNs change

## *Authoritative announcement*

- The IRR has no authorization
- Need to verify that a ASN is authorized to announce a route
- Offline prefix list generation is a 80% improvement
- SBGP and SoBGP build this into the protocol

- No major event to provide a catalyst implementation
- Increased Costs
  - Implementation
  - Provisioning
  - Troubleshooting
- Some customers do not know how to register their routes in the IRR today
  - New solutions add greater complexity
- Low adoption rate
  - Limited number of large ISPs use the IRR to filter customer routes

- Routes take memory
  - Each route takes up space in the RIB (route table,) and the FIB (forwarding table)
  - IPv4 Routing table size is currently 186,594
  - IPv4 Multicast table size is 10,691
  - IPv6 Routing table size is currently 1,513
- L3VPN adds more routes

## *Isolating the Control plane*

- Create a barrier between routing protocols and customer traffic
  - Hide the core
  - Separate the control plane from the data plane
    - Distributed routers
      - Forwarding is separated from Routing
    - MPLS
      - Control traffic, and customer traffic can take different paths
- Easy to determine router loopback addresses
  - Reverse dns hostnames
    - Standard is <interface>.<router>.<pop>.<domain>
    - Means loopback is usually <router>.<pop>.<domain>

## *Options for separation of the control plane*

- no-propagate-ttl
  - BGP Free core
- Isolate the core
  - Ryan McDowell of Sprint presented at NANOG
    - <http://www.nanog.org/mtg-0405/mcdowell.html>
  - Private address space
- L3VPN
- Separate physical network

## *no-propagate-ttl (juniper name)*

- Router does not decrement the ttl of a IP packet within a MPLS LSP
- traceroute cannot see LSRs within the core of the network
  - Shows only ingress, and egress routers
  - Causes operational issue
    - NOC can have a more difficult time troubleshooting
    - Customer cannot perform their own troubleshooting
- Combined with a BGP free core, nodes outside the network cannot see the backbone routers, nor can any return packets be sent from the core back to the node

- Create a dedicated FIB (forwarding table) for Control traffic
  - All control traffic uses one FIB
  - All other traffic uses a different FIB(s)
  - All traffic use the same physical interfaces

- WDM
  - Feasible to have a control channel wave
  - Hardware innovations may help
    - Something like SONET control channel used for ADMs
- Separate circuits
- Difficult to implement in existing implementations

## *Disadvantages of separation of control plane*

- Loss of fate sharing
  - Control plane may not detect when a failure occurs on the data plane
- Requires a separate network just for control traffic
- Increased cost of O&M
  - Different sets of tools to manage and troubleshoot the control plane from the forwarding plane
    - Example: traceroute for control plane, lsptracertool for the forwarding plane
  - More complex
    - Longer to troubleshoot
    - Requires more skilled NOC/ops engineers

## *Management Plane Security*

- Filter
- Open TCP ports
- IPv6
- SNMP
- OOB
- Router Access
  - Remote Access
  - ssh/telnet/ssl
- Passwords
  - Secureid

- Nothing should enter the management plane unless it is permitted
  - SNMP only from specific hosts
  - Ssh only from specific hosts
  - http only from specific hosts
  - TACACS, RADIUS
  - FTP, RCP, TFTP
- Same mechanisms can be used to filter the management plane as the control plane (in Juniper and Cisco)
  - Juniper lo0
  - Cisco Control Plane policing

- Routers listen to TCP/UDP ports
  - Establishing a connection takes resources
  - Attacks have overloaded flooded routers with TCP connections
  - SYN Flood on the router
  - Routers handle connections differently
    - Accept the TCP connection, then check the ACL and close the connection
      - Resources are used for the TCP connection
    - Send ICMP Port unreachable
    - Drop the incoming packet
- Filtering can help

- Do not forget about ipv6
- It is a separate protocol, usually requiring a separate set of filters and security
- Same rules apply as IPv4
  - Filter control plane
  - Filter management plane

- Historically Problematic
  - Buggy
  - Memory leaks
  - Security holes
    - Hard coded communities
- Thought history was behind us, but its not
- Essential Tool for O&M
- Limited security in v1, and v2c
  - V3 has expanded security
  - Lightweight network protocol is part of the appeal
- communities rarely change

- OOB can be built different ways
  - Must be completely separate from the customer network
  - Remote Modem, and console server
  - Separate Frame-relay or ethernet network
- Critical for managing the network
  - Not only for failure but if the network is under attack
    - DOS
    - Owned routers
- Converged networks
  - Frame-relay, ATM over IP
  - VOIP

- Remote access
  - Engineers need access from Conferences in Perth, Australia, and vacations in Hawaii
  - Access must be secure
    - Encrypted
      - Dialup is generally not encrypted
    - Authenticated
  - VPN
  - Management Host
- telnet/ssh
  - telnet and ssh are most common today
  - Ssh should be preferred
    - Ssh key changes
      - Replace the routing engine the key changes
        - Usually just delete the old key

- Need secure passwords
  - Regularly change passwords
  - Use good passwords
- Remote access
  - Conferences, cafés, public networks provide a means for sniffing passwords
- Secureid
  - Provides a one time password
  - Can combine physical security and password security
  - Even if sniffed or keystroke logged, the password is useless after a matter of seconds

## *Data Plane Security*

- DOS
  - Detection
  - Mitigation
- Converged Networks
- QoS

- **Changes in circuit utilization**
  - Traffic graphs work
  - Utilization maps
    - The map will show traffic levels on specific paths not just one interface For example:
      - All paths to Miami are running at 90% utilization
      - San Jose to Dallas to Miami is running at 90% utilization
- **Customer reporting problem**
  - Phone call, email, Website
- **Network analysis tools**
  - Internal tools
  - Commercial tools

All detection methods have a high number of false positives

- Special Events
  - Victoria Secret fashion show, Earthquake
    - CDN is helping, but CDNs can trigger alerts
- slashdot
  - Sudden change in traffic

Customer access during an event

- How does the customer access the website, send email, or call with voip during a DoS attack
  - QoS
- NOC must investigate and notify customer

## Triggers

- Phone Call
- Web Interface
- BGP Community
- Automitigation

## Mitigation strategies

- Blackhole
- ACLs
- Scrubber box

# DoS mitigation Scaling

## Blackhole

- Usually done in the hardware, packets are dropped
- Possible to implement on all routers
- No capacity needed to carry the packet across the network, as the packet is dropped when it enters the network

## ACL scaling

- Hardware limitations
  - Number ACLs
  - Size of ACLs
  - Function of ACLs
- Software limitations
  - Policy in the software may limit the ability to deploy ACLs
- BGP Flowspec may help

## Scrubbing boxes

- Fixed capacity 1G, 2G, 10G (does not matter it has a limit)
  - Same congestion principles as a network
    - Underprovisioned scrubber, DDoS the scrubber, failures
  - Most boxes can handle most DDoS today
    - Some DDoS are greater than the capacity of the box
- Must have enough capacity to carry DDoS traffic from the edge to the scrubber, and then to the customer
- Scrubbers have to hold state many times
  - Creates a limitation
  - DDoS the state in the Scrubber

- Use BGP to specify explicit Network & Transport Layer filters
  - Distributes information about a specific flow beyond src/dst IP address
  - Allows more granular ability to manipulate traffic
    - Filter, rate limti
    - Divert traffic
    - Black hole traffic
    - Policy based routing (PBR)
    - more

## Evolution

- Boxes (1 per network or 1 per pop)
- Clusters (multiple boxes working together)
- ASIC/Linecard
  - New attack comes that the ASIC cannot support, start over

## Common growth curve of Scrubbers

- Single Box
- Multiple boxes
- Regional Boxes (Asia, US, Europe) (Singapore, Hong Kong, Taipei)
- Per Pop boxes
- POP Clusters

# Converged Networks

ntt.net

Control, TDM, Frame-Relay, ATM, Ethernet, IP, L3VPN, Voice, and Video networks are targeted to converge onto a single network

- Most common concern is that IP (public internet,) traffic could cause performance problems with Voice or ATM, etc
- In reality any protocol can interfere with any other
  - Voice traffic on Video traffic
  - Frame-Relay on L3VPN

DOS only happens on the public internet

- ISPs are using l2vpn (ATM/Frame-Relay/Ethernet,) for public internet links
- L3VPNs are used for corporate (and others) networks that are susceptible to viruses and worms

QoS solves this

- QoS adds its own set of problems

The core links need to be **Nx** Larger than the largest edge link

- DoS should only impact the edge
- Customers are demanding 10G links

- Based on Tiered service priority
  - Jitter sensitive traffic is first packet out
    - TDM, Video, Voice
  - Loss sensitive traffic is not dropped, sufficient buffering
    - Packet data, ATM, Frame-Relay, Ethernet
  - IP Traffic uses what is left
    - L3VPN has priority over public IP
  - Control data must be transmitted
- How is it classified
  - Input interface
  - ACL (IP packet)
  - Diffserv/IP Precedence
  - MPLS EXP
- What happens when malicious traffic gets a higher priority
- QoS only impacts traffic during congestion
  - Caused by under provisioning
  - DoS
  - Failure event
- Control traffic set with higher IP precedence bits
  - Places control traffic into a separate queue
  - Default configuration allows all traffic with specific IP precedence bits to be placed into same queue as the control traffic
  - Reset IP precedence bits or change classification on the edge

- Limitations
- Vulnerabilities
- O&M

## *Hardware Limitations*

- Pps performance
- Bps performance
- Flow performance
- Feature performance
  - ACLs
  - Accounting
  - Recirculating the packet
  - Punting the packet

- Software has bugs, Hardware has bugs
  - ASICs are becoming much more complex, leading to a higher likelihood of bugs
  - Hardware has a longer lead time to repairs
    - Usually measured in years
    - Software sometimes can work around the problem

- Must be able to manage the hardware
- Counters
  - Queue
  - Interface
  - Lsp
  - MAC Address counters
- Flow data
  - Source/Destination, IP/Port, BGP next-hop, input/output interface, TCP flags
- Traceback
  - Input interface, and input LSP, SRC MAC address

- NOC

### Top 5 security events handled by our NOC

- phishing
- bots
- infected hosts
- spam
- Ddos
  - Reduced number of impacting DDOS recently

## *Most useful tools used by our NOC*

Netflow data

Darknet

Router stats

Network stats

Security lists

- nsp-sec
- bugtraq
- cert list

Abuse email ticketing system

## More refined Netflow stats

- Commercial tools
- Further internal development

## Deep packet inspection

- How do we look into the packet beyond our ability today
- How far do we need to look into the packet (100%?)

- Security is a big complex problem
- Always changing
- Preparation must be made ahead of time
- Many pieces that creates complexity
- Information sharing is critical



*Thanks*

ntt.net

Peter Schoenmaker  
pds@ntt.net