



NETWORKERS 2004

ADVANCED PREFERENTIAL IP AND TELEPHONY SERVICES IN THE INTERNET

SESSION VVT-4002

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

1

Lessons Learned 11 September 2001

Cisco.com

- **There were many simultaneous individual failures**
 - Some could have been avoided by using better procedures
 - Many were automatically routed around or recovered from
- **On the whole, the Internet is what worked that day**
 - The Internet is reliable because it is distributed in control and operation
 - Even if individual parts perform in new and exciting ways
- **The Internet is useful for emergency services such as police, fire, and medicine, and for military use**
- **The Internet is now critical national communication infrastructure**

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

2

Issues We Would Like to Cover

Cisco.com

- **Existing Public Switched Telephone Network (PSTN) Voice Services**
 - MultiLevel Precedence and Preemption (MLPP)
 - Government Emergency Telecommunications Service (GETS)
- **Extensions to the Internet**
 - Application to Voice, Video, and Real-Time Services
 - Application to File Transfer, Email, and Messaging
- **Key Concept:**
 - PREFERENTIAL SERVICE OFFERED TO AN AUTHENTICATED USER**
- **This Is Not e-911**
 - See Session VVT-4000

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

3

MultiLevel Precedence and Preemption (MLPP)

Cisco.com

In Military Communications, a Priority Scheme

- For assigning one of **several precedence levels** to specific calls or messages so that the system handles them in a predetermined order and time frame
- For gaining **controlled access to network resources** in which calls and messages can be preempted only by higher priority calls and messages
- That is recognized only within a predefined domain, and
- In which the precedence level of a call outside the predefined domain is usually not recognized

<http://www.its.bldrdoc.gov/fs-1037/dir-023/3416.htm>

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

4

Government Emergency Telecommunications Service (GETS)

Cisco.com

- **Civilian counterpart to MLPP**
 - Preemption is not permitted by FCC
- **Mechanisms used:**
 - Reservation of last trunk on a multiplexed line
 - Queuing for next available trunk
 - Expanded call routing options



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

5

Key Points from a QoS Perspective

Cisco.com

- **Many links along a path are very high speed and offer very high bandwidth**
 - LANs, Optical WANs
- **Other links along a path run at dramatically lower speed**
 - Throughput is constricted at these points:
 - Interconnections between networks
 - Forward deployments, especially satellite, line-of-sight, and mobile
 - “Last mile”
- **QoS issues are found in constriction points, not LANs and optical backbones**

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

6

Issues That Have to Be Addressed

Cisco.com

Issues at Several Layers

- **Politics:** Who gets to use preferential access?
- **Application layer:** is there a difference in queuing or other resources?
- **Session Layer:** do we let script kiddies use this? How do we keep them out?
- **Can a transport provide predictable throughput?**
- **Can the Network layer provide guarantees to the packets?**

Political
Religion
Application
Presentation
Session
Transport
Network
Link
Physical

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

7

RFC 1633: Types of Applications

Cisco.com

Two Broad Classes:

- **Those with fixed requirements of the network**
“Real Time” applications include voice, video, telemetry, and any other application with an **IRREDUCIBLE LOWER BOUND** on its bandwidth requirements or an upper bound on the jitter and loss it will accept
- **Those that will adapt to a wide range of network environments**
“**ELASTIC**” **APPLICATIONS ADAPT** themselves to any bandwidth as long as it is consistently available; Examples include (but are not limited to) TCP and SCTP-based applications

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

8

Messages from DARPA* and NCS**

Cisco.com

- **Statistical services like the Internet provide statistical guarantees**
- **Sometimes, people need non-statistical guarantees:**

**A certain amount of bandwidth between here and there,
no excuses**

Move a file across the network in a fixed amount of time

**Deliver a message to a set of people within a fixed period
of time**

* **DARPA: Defense Advanced Research Projects Agency**

****NCS: Department of Homeland Defense, National Communication Service**

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

9

PREFERENTIAL REAL-TIME (VOICE/VIDEO) SERVICES



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

10

Important Operational Characteristics

Cisco.com

- **Need to be able to authenticate and authorize certain calls**
- **MLPP: Need to be able to preempt calls**
 - At handset—incoming call preempts existing call
 - In network—new bandwidth requirement preempts lower precedence usage of the bandwidth
- **Need to signal to users using standard signals**
 - Chime/tone indicating preemption
- **Need to preserve existing calls at all precedences**

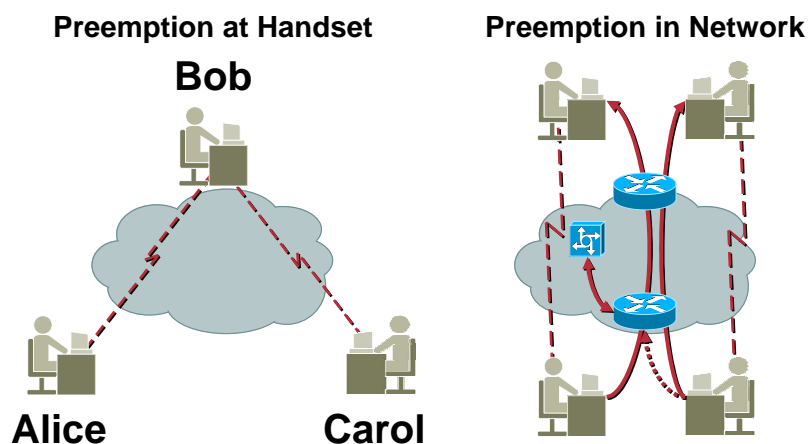
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

11

Two Fundamental Preemption Cases

Cisco.com



VVT-4002
9840_05_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

12

End System Preemption

Cisco.com

- Deals with case where call with elevated precedence is placed to a handset that is in use
 - Alice calls Bob who is talking with Carol at lower precedence
- H.323:
 - H.460.mlpp
- SIP Solution:
 - SIP Resource Priority Header
 - Label call with precedence level
 - SIP Call Failure Reason Code
 - Reason = "Call Preempted"



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

13

Network Admission/Preemption: Call Admission Control (CAC) Solutions

Cisco.com

- Heuristics based on loss
 - Differential loss among different calls
 - Calls experiencing high loss might self-drop
- Server-based call admission
 - Call counting
- Network bandwidth admission
 - Using signaling with server CAC

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

14

Issues with Heuristics Based on Loss

Cisco.com

- In a word, not deterministic
- In overload case, all calls at a precedence see roughly equal loss
 - All calls at a level hear a lot of noise
 - Users tend to re-call "due to bad circuit"
- If instruments apply algorithms to measure and respond to loss by dropping
 - Result is indeterminate
 - May drop too many calls
 - May incorrectly present signals to users

VVT-4002
9840_05_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

15

Loss and Jitter Are Issues...


Cisco.com

- ...because jitter and loss directly affect signal quality



ON THE PERFORMANCE OF ILBC OVER NETWORKS WITH BURSTY PACKET LOSS
Mingyu Chen and Manohar N. Murthy

- In the end:
 - Service preserves higher precedence calls
 - All calls that are not preempted should enjoy reasonable sound quality



Packet Loss and Packet Loss Concealment
Nortel Networks

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

16

Server-Based Call Admission

Cisco.com

- **Server-based admission implies that server knows the network**
 - It has knowledge of its own domain and neighboring domains, and policies regarding codecs and permitted number of simultaneous calls
 - It does not know routing or ambient traffic loads
- **Servers are often located away from tactical deployments and specific offices**
 - And specifically not in routers

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

17

Network Bandwidth Admission

Cisco.com

- **It is possible, using RSVP to**
 - Use control plane signaling to deterministically authorize/preempt bandwidth (RFC 2205)
 - Enable SIP to start data transmission only when network bandwidth has been authorized (RFC 3312)
 - Aggregate calls within backbones (RFC 3175)
 - Limit state to systems with bottleneck links
- **NSIS is a possible future**
 - IETF "Next Steps in Signaling" WG; work in progress

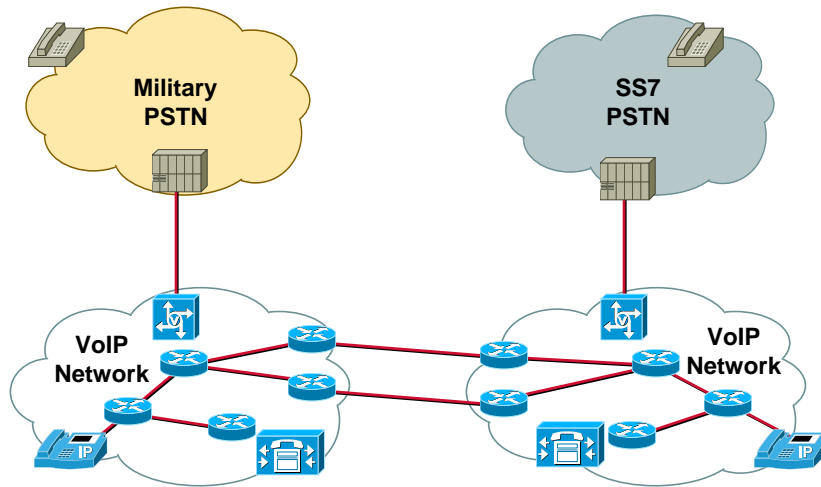
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

18

Control Paths May Not Follow Data Paths

Cisco.com



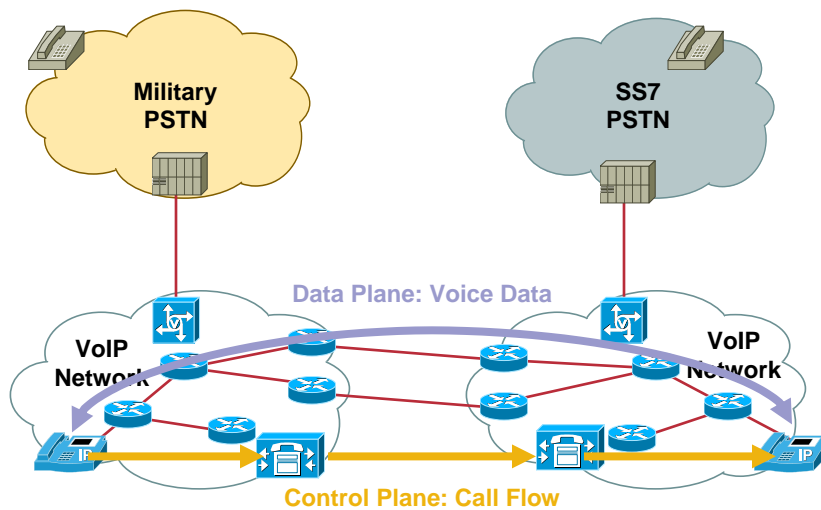
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

19

VoIP Call Control and Data

Cisco.com



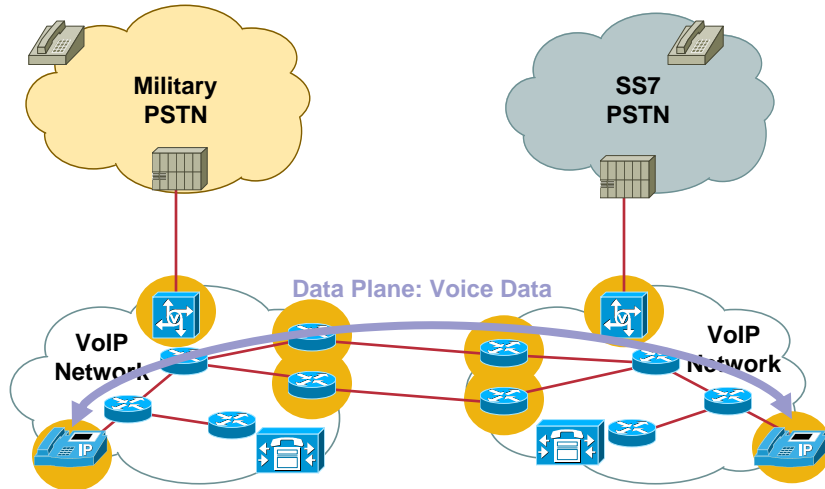
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

20

Where to Configure Signaling

Cisco.com



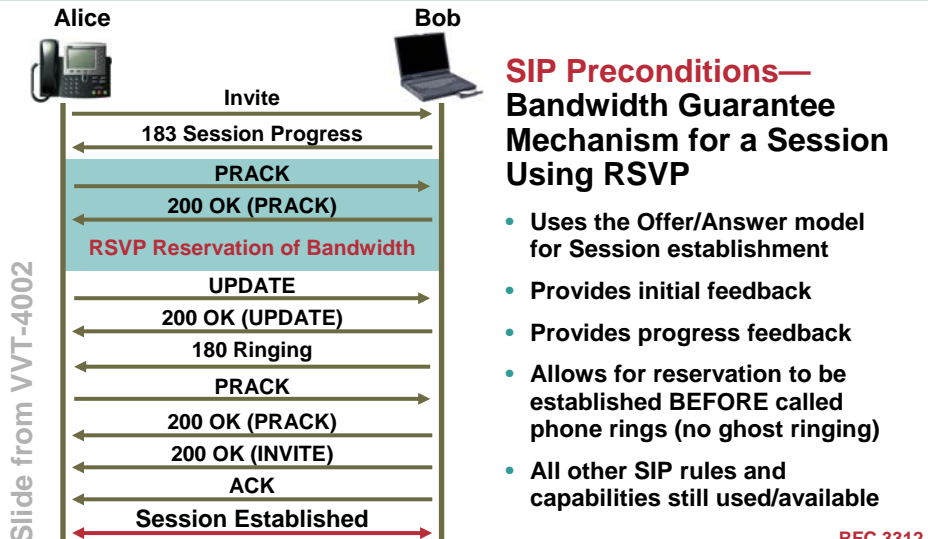
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

21

SIP (QoS) Preconditions

Cisco.com



SIP Preconditions— Bandwidth Guarantee Mechanism for a Session Using RSVP

- Uses the Offer/Answer model for Session establishment
- Provides initial feedback
- Provides progress feedback
- Allows for reservation to be established BEFORE called phone rings (no ghost ringing)
- All other SIP rules and capabilities still used/available

RFC 3312

Slide from VVT-4002

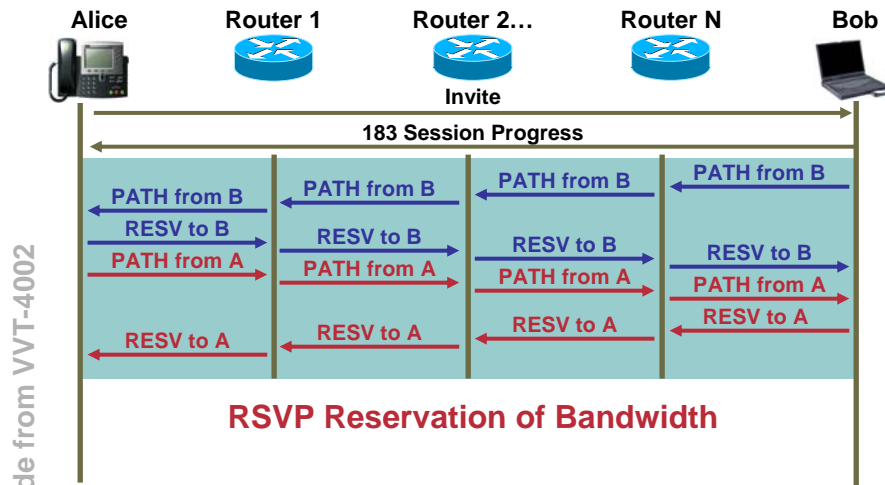
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

22

SIP (QoS) Preconditions

Cisco.com



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

23

Session AAA Procedures in RSVP

Cisco.com

- **RSVP session types**
 - IPv4/IPv6: RFC 2205 {IP addresses, IP Protocol, UDP/TCP Ports}
 - IPSec: RFC 2207 {IP addresses, IPSec Association}
 - Aggregated: RFC 3175 {IP Addresses, DSCP}
- **Policy Data Object**
 - RFC 2750 RSVP Extensions for Policy Control
 - RFC 2753 A Framework for Policy-based Admission Control
 - RFC 2872 Application and Sub Application Identity Policy Element for Use with RSVP
 - RFC 3182 Identity Representation for RSVP
- **Policy data object identifies:**
 - Requestor, by personal identity or some aggregate identity
 - Priority of reservation as a numeric value

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

24

Proposal

Cisco.com

- **Use RSVP signaling in control plane**
 - RFC 2205, 2207, or 3175
 - Primarily configured at interchange points
 - Only on systems with bottleneck interfaces
 - With AAA/policy/preemption capabilities as described in RFCs
- **Use EF PHB for VoIP traffic data plane**
 - RFC 3246/3247
 - Flag voice with DSCP, apply priority and simple policing to sum of bandwidths of calls

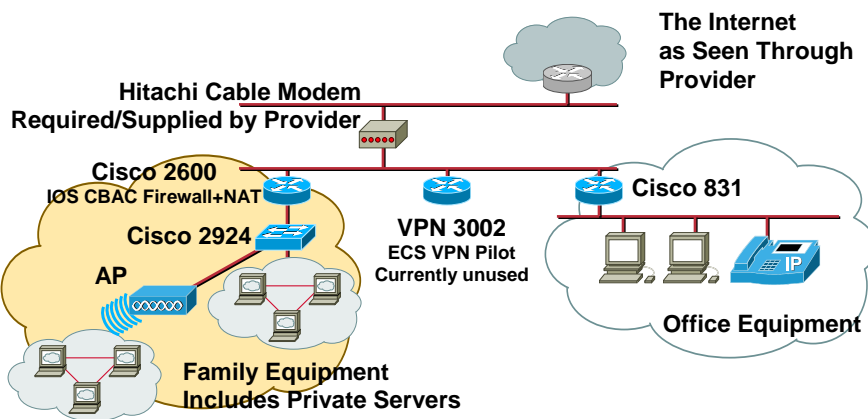
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

25

Home Network Diagram

Cisco.com



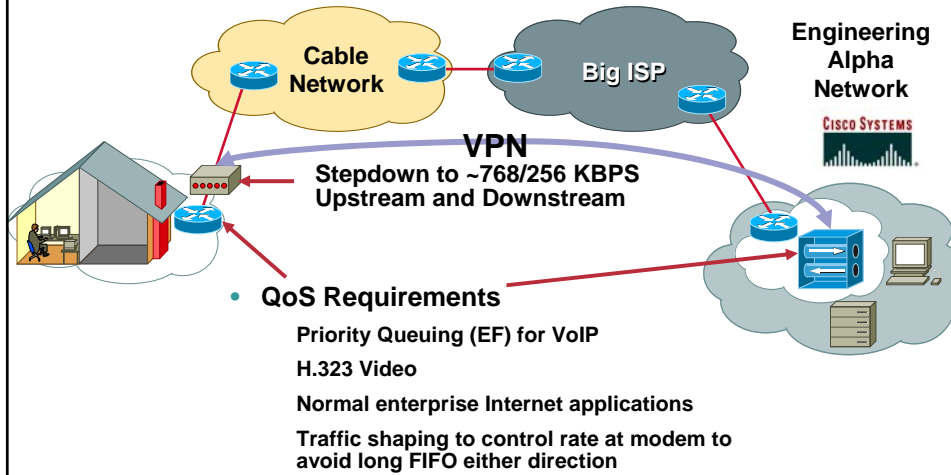
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

26

Big Picture Network Diagram

Cisco.com



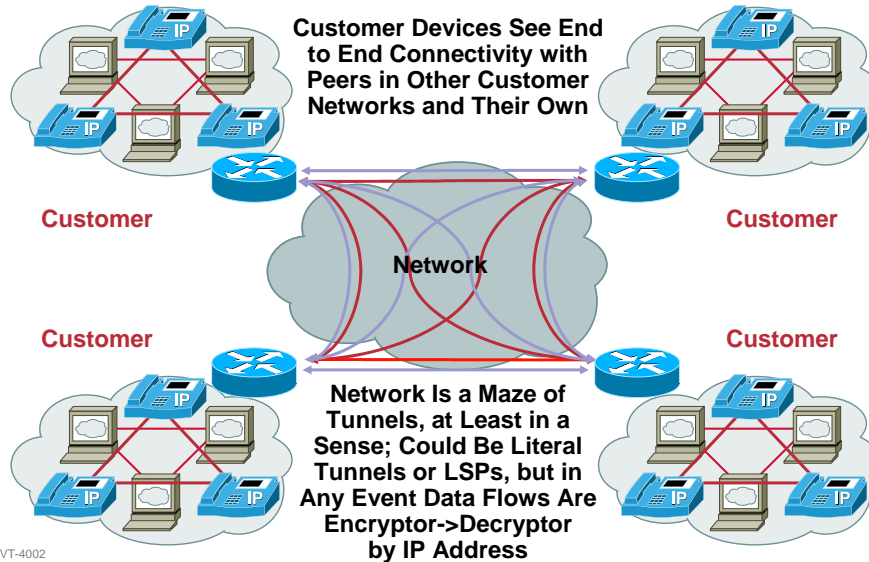
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

27

Mental Model of a VPN Network

Cisco.com



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

28

Tunnels and Aggregation

Cisco.com

- **Kinds of aggregates**
 - MPLS
 - GRE, IP/IP, IPSec firewall traversing
 - IP traffic traversing the same ingress and egress
- **Ways to aggregate calls**
 - RFC 2207 RSVP/IPSec for IPSec VPN
 - RFC 3175 Aggregated RSVP for others

VVT-4002
9840_05_2004_c2


© 2004 Cisco Systems, Inc. All rights reserved.

29

Structure of VPN Edge Router

Cisco.com

- **Customer Message**
 - Received from customer Network
 - Routed to network
 - Encrypted, adding IPSec ESP
 - Encapsulated in network IP header
 - Routed in tunnel
- **Network Message**
 - Received in tunnel
 - Decrypted
 - Routed, usually to customer, but also possible to route to another tunnel
- **Example:**
 - Configuration in VPN Network
 - Tunnels set up via Routing
 - Keys exchanged by combination of configuration and key exchange under "Unity" security architecture
 - Encrypt/decrypt performed in box labeled "security association" using relevant key



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

30

RFC 2207: RSVP for IPsec Data Flows

Cisco.com

- **Enables RSVP to signal QoS Flow Specification per Security Association per IP source/destination address pair**

Two systems may have multiple security associations, and therefore SPIs

Data may be identified with same DSCP, as this is in the control plane

- **Flow characteristics of a tunnel are the sum of the flow characteristics of the traffic in the tunnel**

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

31

RFC 3175: RSVP Aggregation

Cisco.com

- **Designed to permit aggregation of reservations**

Service Provider Voice...

- **Ways to identify affected traffic**


IP Source/Destination/DSCP

Tunnels and LSPs

- **Mechanism:**

RSVP reservation for aggregate is the sum of the reservations using it

Traffic in aggregate identified/services by DSCP



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

32

Preemption within an Aggregate

Cisco.com

- **Procedure for preemption within VPN core**

New call arrives at advanced priority

May result in new aggregate reservation for one call or increase in bandwidth of existing reservation

RESV_ERROR message along routine aggregate

Instructs de-aggregator to reduce bandwidth


MLPP policy used to decide which customer call to drop

Lower priority aggregate

Reinstated at lower bandwidth or removed

- **Key implication:**

Customer and network sides of VPN router have to exchange information



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

33

Key Information Leakage Observation:

Cisco.com

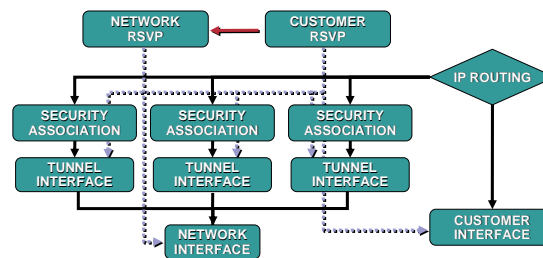
Traffic in the network can be measured by any device in the network

Meta-statements about that traffic, such as reservation information or MPLS-TE bandwidth commentary, conveys no information that could not be obtained by measurement

Proposal

Cisco.com

- **Customer:**
Same as in normal network as far as it knows, it is a normal network
- **Network:**
Advertise PATH/RESV for sum of sessions using a tunnel
If new network reservation refused, apply MLPP policy to refuse or preempt customer reservations
If old network reservation preempted, repeatedly apply MLPP policy to customer reservations until network reservation is accepted
- **Use RSVP 2207/3175**
Reservations for IPSec security associations, one SA for best effort traffic and one for each level of VoIP traffic
Bandwidth of a tunnel \leq sum of mean rates of sessions using tunnel



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

35

Living with Mobility

Cisco.com

- In a mobile ad hoc network, concerns have been raised that reservations may not be in place when something moves
- **Scenarios:**
 - Pure differentiated services: no limit on bandwidth usage
 - Location-based CAC: need location sensitivity at each aggregation router
 - Circuit models: lose traffic (tear down or repair circuit)
 - Signaled QoS: IP traffic switches before reservation, but reservation follows quickly
- **From my perspective,**
 - None are perfect, but
 - Signaling is less painful than circuit switch implications and
 - Less operational issue than putting a call manager in every backpack

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

36

Issues Addressed in This Solution

Cisco.com

- **Who gets authenticated to use preferential access?**

That is a policy question, not addressed in the technology

- **Application layer issues:**

Endpoint preemption, Authentication, Authorization in SIP

- **Session Layer controls:**

RSVP provides bandwidth/route-aware CAC

Session AAA issues handled in policy

- **Network layer guarantees?**

Yes, now that we know we have available bandwidth

Political
Religion
Application
Presentation
Session
Transport
Network
Link
Physical

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

37

PREFERENTIAL ELASTIC SERVICES



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

38

Preferential Services for Elastic Applications

Cisco.com

- **Example applications**
 - File transfer
 - Transaction applications
 - Instant messaging
 - Electronic mail
- **Issue: Predictable access during congestive collapse**
 - Large proportion of network out of service
 - Ongoing electronic attacks
 - Heavy Internet traffic from users
 - Emergency management needs access

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

39

Issues That Have to Be Addressed

Cisco.com

Issues at Several Layers

- Who gets authenticated to use preferential access?
- Application layer: Is there a difference in queuing or other resources?
- Session layer: Do we let script kiddies use this? How do we keep them out?
- Can a transport provide predictable throughput?
- Can the network layer provide guarantees to the packets?

Political
Religion
Application
Presentation
Session
Transport
Network
Link
Physical


VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

40

Network Layer Procedures

Cisco.com



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

41

Two Approaches to Making Guarantees to IP Traffic

Cisco.com

- **Drop precedence**
 - Permit traffic with multiple DSCPs to use the same queue
 - Make some DSCPs compete more effectively for queue bandwidth
- **Class provision**
 - Divide traffic with differing DSCPs into different queues
 - Give different guarantees—or no guarantees—to various queues

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

42

Drop Precedence Effects

Cisco.com

- **RFC 2581 TCPs respond to loss by slowing down**
Therefore, loss is a signal to a sending TCP
- **RFC 2597 Assured Forwarding**
Assures at least minimal service to each competing session
“Excess” users tend to absorb drops, and therefore slow down
Tends to guarantee fairness
- **Assigned drop precedence (“routine”, etc.)**
Preferred sessions still compete with lower priority sessions
No guarantees to preferred sessions beyond “less loss”

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

43

Effect of a Class Definition

Cisco.com

- **Similar to TDM definition of a multiplexing channel:**
Guarantees a minimum level of bandwidth if there is traffic to use it
- **But “work conserving”:**
Unused bandwidth is redistributed among other classes
- **Implication:**
Cost of an unused class is nil
Makes bandwidth guarantee to traffic should traffic show up for it
Radical effect on traffic: traffic in another class that is using redistributed bandwidth suddenly reduced by amount that preferred class uses

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

44

Building Preferential Classes

Cisco.com

- **Services defined:**
 - GETS requires one “emergency traffic” class
 - MLPP requires several graded traffic classes, each for all elastic traffic at that class
- **My suggestion:**
 - Define classes and apply bandwidth to them
- **Issue in this:**
 - How do we ensure that only authorized users use these classes?

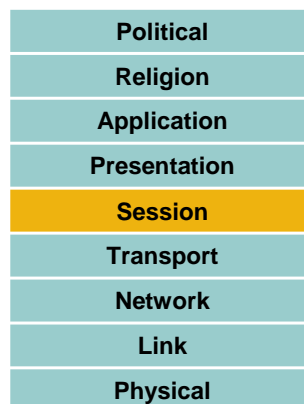
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

45

Security: Who Is Authorized to Get Preferential Access?

Cisco.com



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

46

Identification, Authentication, Authorization

Cisco.com

- **GETS/MLPP Policy:**
 - Provides service to an **AUTHENTICATED USER** who **REQUESTS** the service
- **Need to provide ways to**
 - Authenticate the user
 - Let him/her request the service
 - Authorize the usage
 - Provide access to the authenticated user

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

47

Authenticate the User

Cisco.com

- **This uses an existing AAA function:**
 - TACACS, RADIUS, Diameter, etc.
 - Uses standard AAA facilities and procedures
- **Requirements:**
 - Simple/scalable to deploy to hundreds of thousands of users
 - Strong enough to prevent unauthorized access
 - Examples: one-time password, cryptographic authentication with PKI, Kerberos

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

48

Authorize the Usage

Cisco.com

- On presentation of request and credential, must be simple to implement (such as updating ACL)
- Effect is to provide privileged access to the authenticated user

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

49

Let Him/Her Request the Service

Cisco.com

- Request could be accomplished with
RSVP for a TCP session + DSCP, or
Bandwidth broker* that modifies ACL at ingress
node + DSCP
Here there be dragons
- Key requirements:
It must scale to tens to hundreds of simultaneous requests
Must be easily issued by existing applications

* Bandwidth Broker: A Central Routing/Reservation Management System


VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

50

File Transfer Services

Cisco.com



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.


51

File Transfer Is Particularly Part of This

Cisco.com

Drivers for Current Research Agenda

- **Military requirements for predictability**
Large delay*bandwidth and large error rates
- **Research requirements for large file transfer**
Large delay*bandwidth
- **VPNs on the Internet**
Maximizing serviceability on the backbone




VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

52

What Users Would Like...

- “If I have stated capacity between here and there, could I please move a file at that rate?”
- Users tend to respond in surprise when they cannot, or to assert that Internet technology is deficient

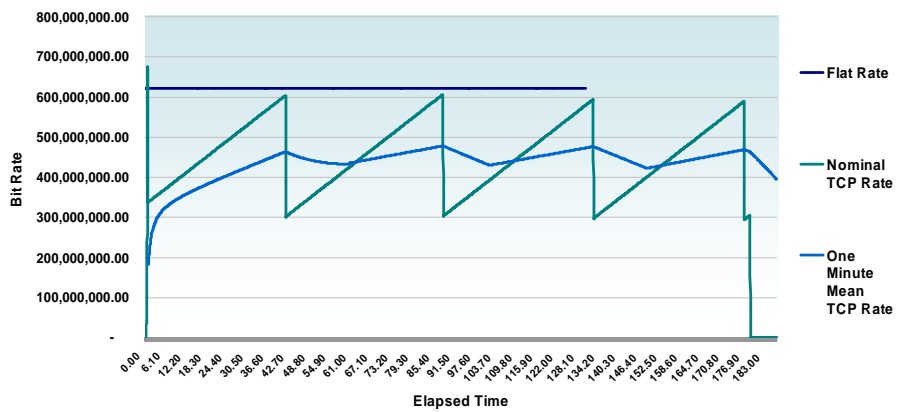


VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

TCP Actual Performance

Performance Curves



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

Australian Defense Solution

Cisco.com

Simple File Transfer:

- Divide file into enumerated blocks
- Place blocks into transmission set
- While transmission set is not empty
 - Send enumerated blocks to receiver
- Receiver acknowledgements remove blocks from transmission set on receipt at sender

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

55

Behavior of That Solution in Time



Cisco.com

- Delivers packets in time proportional to $(1+\epsilon)^n$

ϵ is the error/loss rate

n is the number of phases

- Roughly equal to transmission rate at bottleneck less errored frames



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

56

Calculating the Maximum Rate at Which TCP Sends

Cisco.com

- **RFC 2581 (Reno and NewReno) “standard” TCP**
Heuristics adjust window against loss

$$cwnd_{i+1} = \begin{cases} \text{loss: } cwnd_i/2 \\ \text{no loss: } cwnd_i + mss \end{cases}$$

$$rate = \frac{\text{effective window}}{\text{round trip time}}$$

- **CalTech FAST TCP**
Heuristic adjusts window against varying delay

$$cwnd_{i+1} = cwnd_i \frac{\text{least RTT}}{\text{mean RTT}} + mss$$

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

57


Issues and Approaches in TCP or SCTP Performance: Measuring Available Capacity

Cisco.com

- Can we generate a filtered mean rate rather than playing with heuristics applying to the window?
- Ack pairs constitute packet pairs

$$\text{mean rate} = \text{mean} \left(\frac{\text{amount acknowledged}}{\text{acknowledgement interval}} \right)$$

- Treat this as an upper bound rate
Maximum number of octets sent during relatively small interval shaping traffic to minimize queues
- How does rate-control management interact with “friendliness”?



Source: Keshav, Packet-Pair Flow Control
<http://www.cs.cornel.edu/skeshav/doc/94/2-17.ps>

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

58

Upper Bound Rate

Cisco.com

- We can therefore calculate the rate at which to send

$$\text{mean available rate} = \text{moving average} \left(\frac{\text{amount acknowledged}}{\text{acknowledgement interval}} \right)$$

- And as a result calculate the window size to use instead of estimating it

$$\text{window upper bound} = \text{mean available rate} * \text{RTT}$$

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

59

Issues and Approaches in TCP Performance: Acknowledgement, Loss, and Retransmission

Cisco.com

- Use rate measurement for congestion avoidance
- Retransmit on loss

Detect loss by timeout: mean RTT + mean DRTT

- SACK with windows $O(N * \text{delay} * \text{bandwidth})$ more effective

Redefine “effective window” as mean RTT * mean available rate, and keep that amount outstanding

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

60

We Should Therefore...

Cisco.com

Be Able to Implement TCP in a Manner That:

- Delivers traffic at a predictable rate
- Is compatible with current generation TCPs
- Interoperates well with those TCPs


VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

61

Issues in the Application: Electronic Mail as an Example

Cisco.com



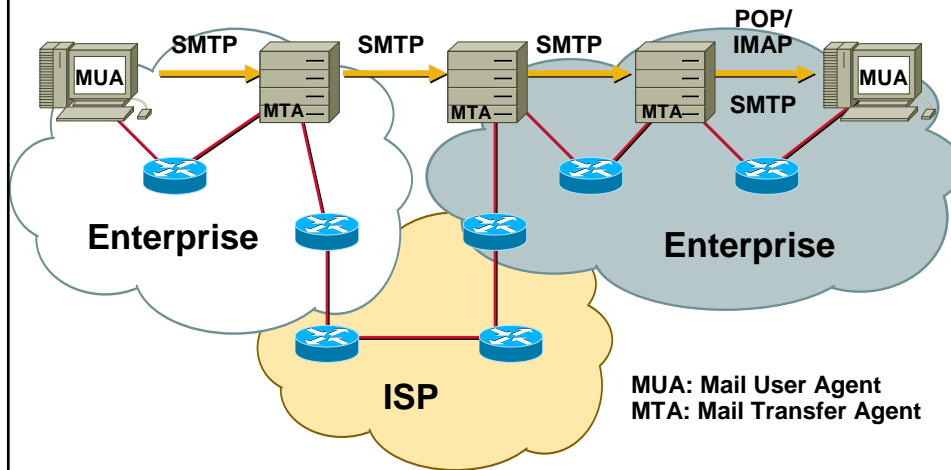
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

62

Electronic Mail as a Privileged Service

Cisco.com



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

63

Issues to Consider

Cisco.com

- **In MTA->final MUA**
 - Need a “push” protocol, not a “pull” protocol
 - Need notification to user to read priority message
- **In initial MUA and each MTA**
 - Need recognition of priority messages
 - Changes in policy:
 - Do not use exponential backoff on failure
 - Quick send/resend
 - Quick “bounce” on failure
- **Is this for everyone?**
 - Probably just for authenticated users: SMTPAUTH
 - Probably just when requested: SMTP Priority

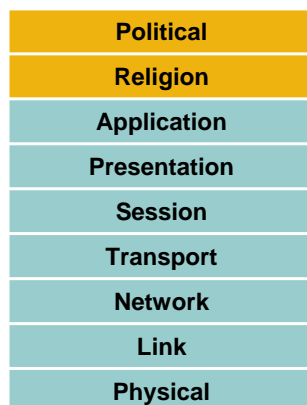
VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

64

Now: Who Do You Give Access to?

Cisco.com



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

65

“REVOLUTIONS BEGIN IN THE COUNTRYSIDE...”

MAO TSE TUNG

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

66

BUILDING PREFERENTIAL SERVICES



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.


67

Key Things to Note

Cisco.com

Building a Preferential Service Requires Technology in Many Layers

- **Application**
Identification and implementation issues in services
Choice of protocols and features
- **Session**
Strong authentication of individual users or functional groups of users
- **Transport**
Appropriate procedures in end system implementations
- **Network**
Application of bandwidth to the right sessions according to an authenticated policy



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

68

“Everything is very simple in war, but the simplest thing is difficult. These difficulties accumulate and produce a friction, which no man can imagine exactly who has not seen war.

So in war, through the influence of an ‘infinity of petty circumstances’ which cannot properly be described on paper, things disappoint us and we fall short of the mark.”

CARL VON CLAUSEWITZ
ON WAR, CHAPTER VII, “FRICTION IN WAR”



Preferential Services: Not Hard, Not Trivial

Cisco.com

- **Technology is mostly in place, but not in use for this purpose**
 - Obvious next step: harness it to the purpose
 - Many specific recommendations
- **We must approach the problem with our eyes open**
 - Careful service definition at each layer
 - Consideration of the specific needs of users one wishes to give preference to

Complete Your Online Session Evaluation!

Cisco.com

- WHAT:** Complete an online session evaluation and your name will be entered into a daily drawing
- WHY:** Win fabulous prizes! Give us your feedback!
- WHERE:** Go to the Internet stations located throughout the Convention Center
- HOW:** Winners will be posted on the onsite Networkers Website; four winners per day

And now a word from the sponsor...

VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

71

CISCO SYSTEMS



VVT-4002
9840_05_2004_c2

© 2004 Cisco Systems, Inc. All rights reserved.

72