Network Positioning System

Pranav Dharwadkar

Product Manager, Cisco,

(pranavd@cisco.com)

September 2011

Picture of the Problem Space

Clients don't have a view of the underlying network topology

DNS employs either a round robin or random scheme

If B is selected, then content download crosses expensive transit links costing ISP $$$ dollars

If A is selected, many hops away, results in bad user experience

In both cases client performance could be sub-optimal

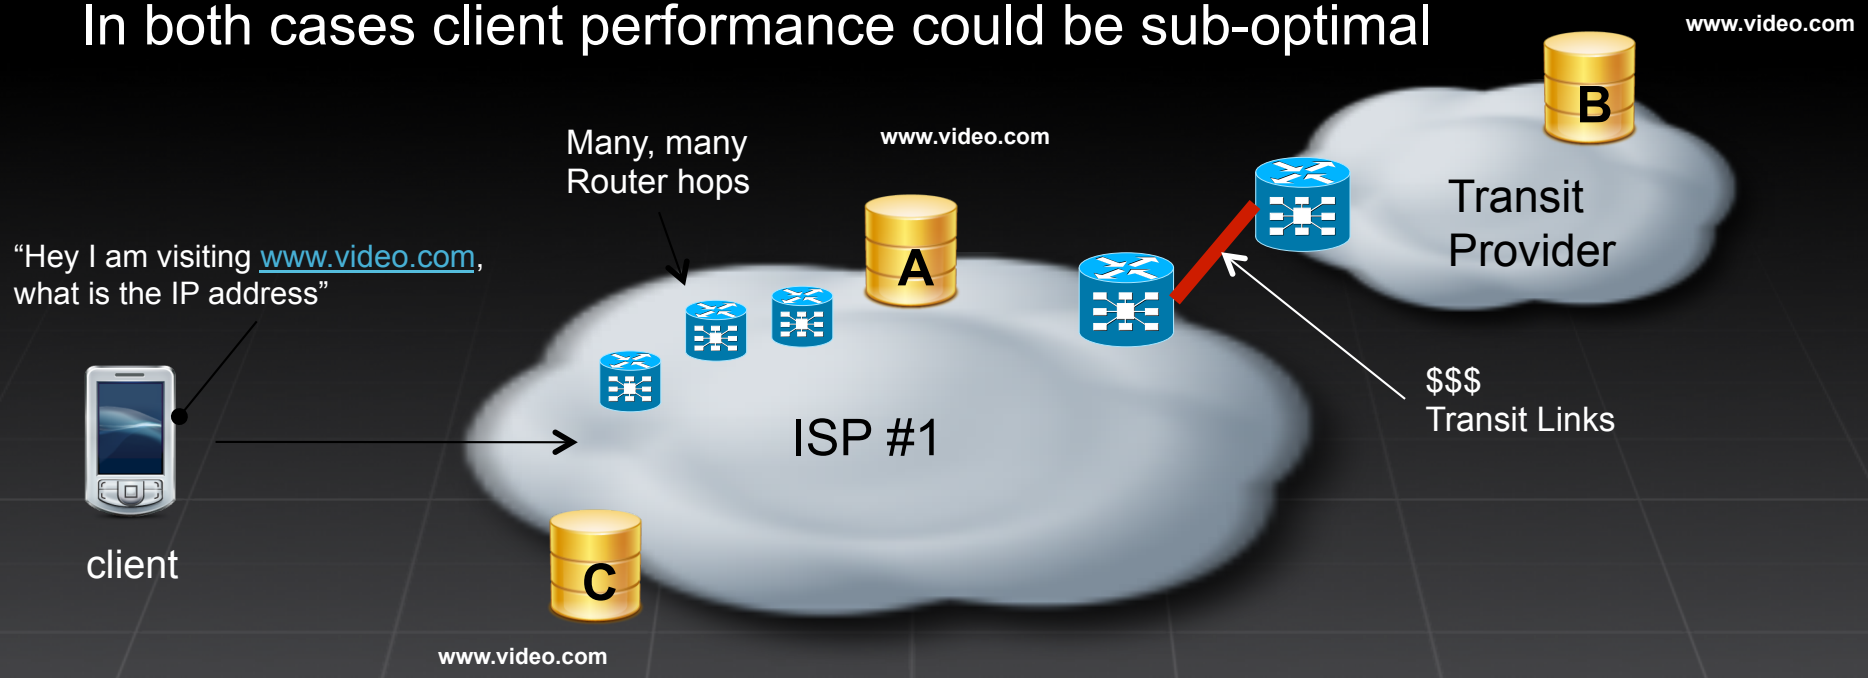

A Better Solution would be …

Application consults the underlying IP topology to determine which host among the several alternatives is “closest” to the client.

Client uses this information and selects C

Application – Topology Interaction is called Proximity

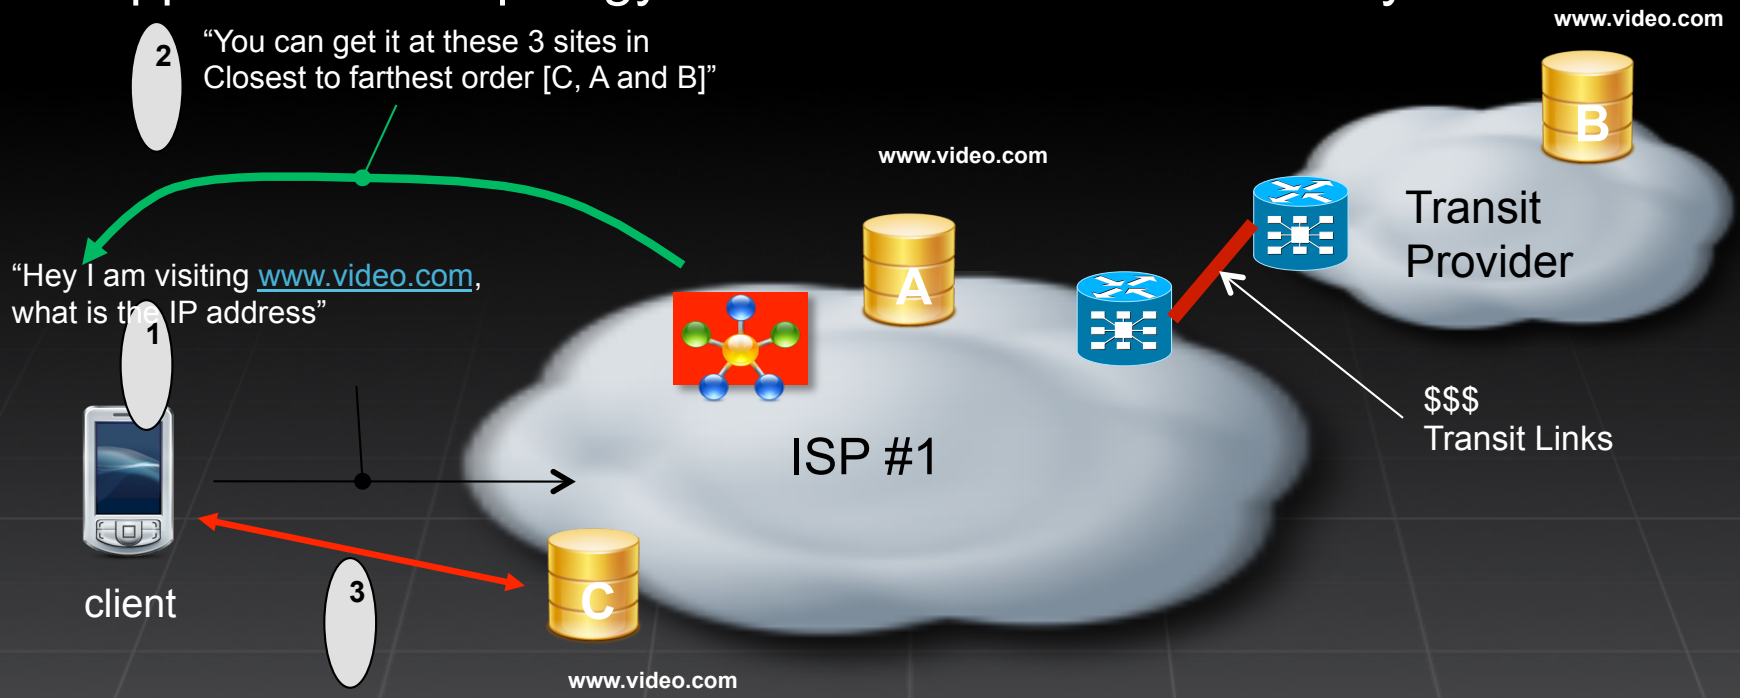

Why Performance Matters

bing found that a **2** second **slowdown** = **4.3%** **reduction** in revenue/user

Google stated that a **400**ms **delay** → **0.59%** **fewer** searches/user

shopzilla reduced page load times from ~7 to ~2 seconds, leading to a

7–12% **increase in revenue** AND **50%** **reduction in operating costs**

AOL noticed that users that experienced fastest page load times view **50%** **more pages per visit** than users experiencing the slowest page load times

Cable&Wireless Worldwide

Compuware

Enabling Technology: ALTO, NPS

- ALTO (defined in the IETF)
  - Application Layer Traffic Optimization (ALTO) defines an interface through which an application can request guidance from the network, e.g. which can be used for service location or placement
    - No need to know atomic topology details
    - Need to preserve confidentiality between layers
  - ALTO does not define the mechanisms used for deriving network topology/infrastructure information or preference
- NPS
  - Network positioning system (NPS) is a specific implementation of mechanisms and algorithms leveraging routing and IP/MPLS infrastructure layer database (such as ALTO), performance, and policy information

NPS: Query/Reply

- The proximity query in the most general form is

  My source address is 192.168.10.1, tell me
  which is the "Best" destination to choose from amongst
  192.168.20.44
  192.168.43.32
  192.168.65.76
  192.168.32.21

- "Best" could be

  Closest (IGP metrics, TE Tunnel, BGP Med)

  Policy (premium users choose dest 1, 2;
  non-premium users choose dest 3, 4)

  Best Link Utilization (choose cheaper link until BW exceeded)

  Time of Day Policy (during day use west coast link,
  during night use east coast link)

NPS Architecture – Layer Separation

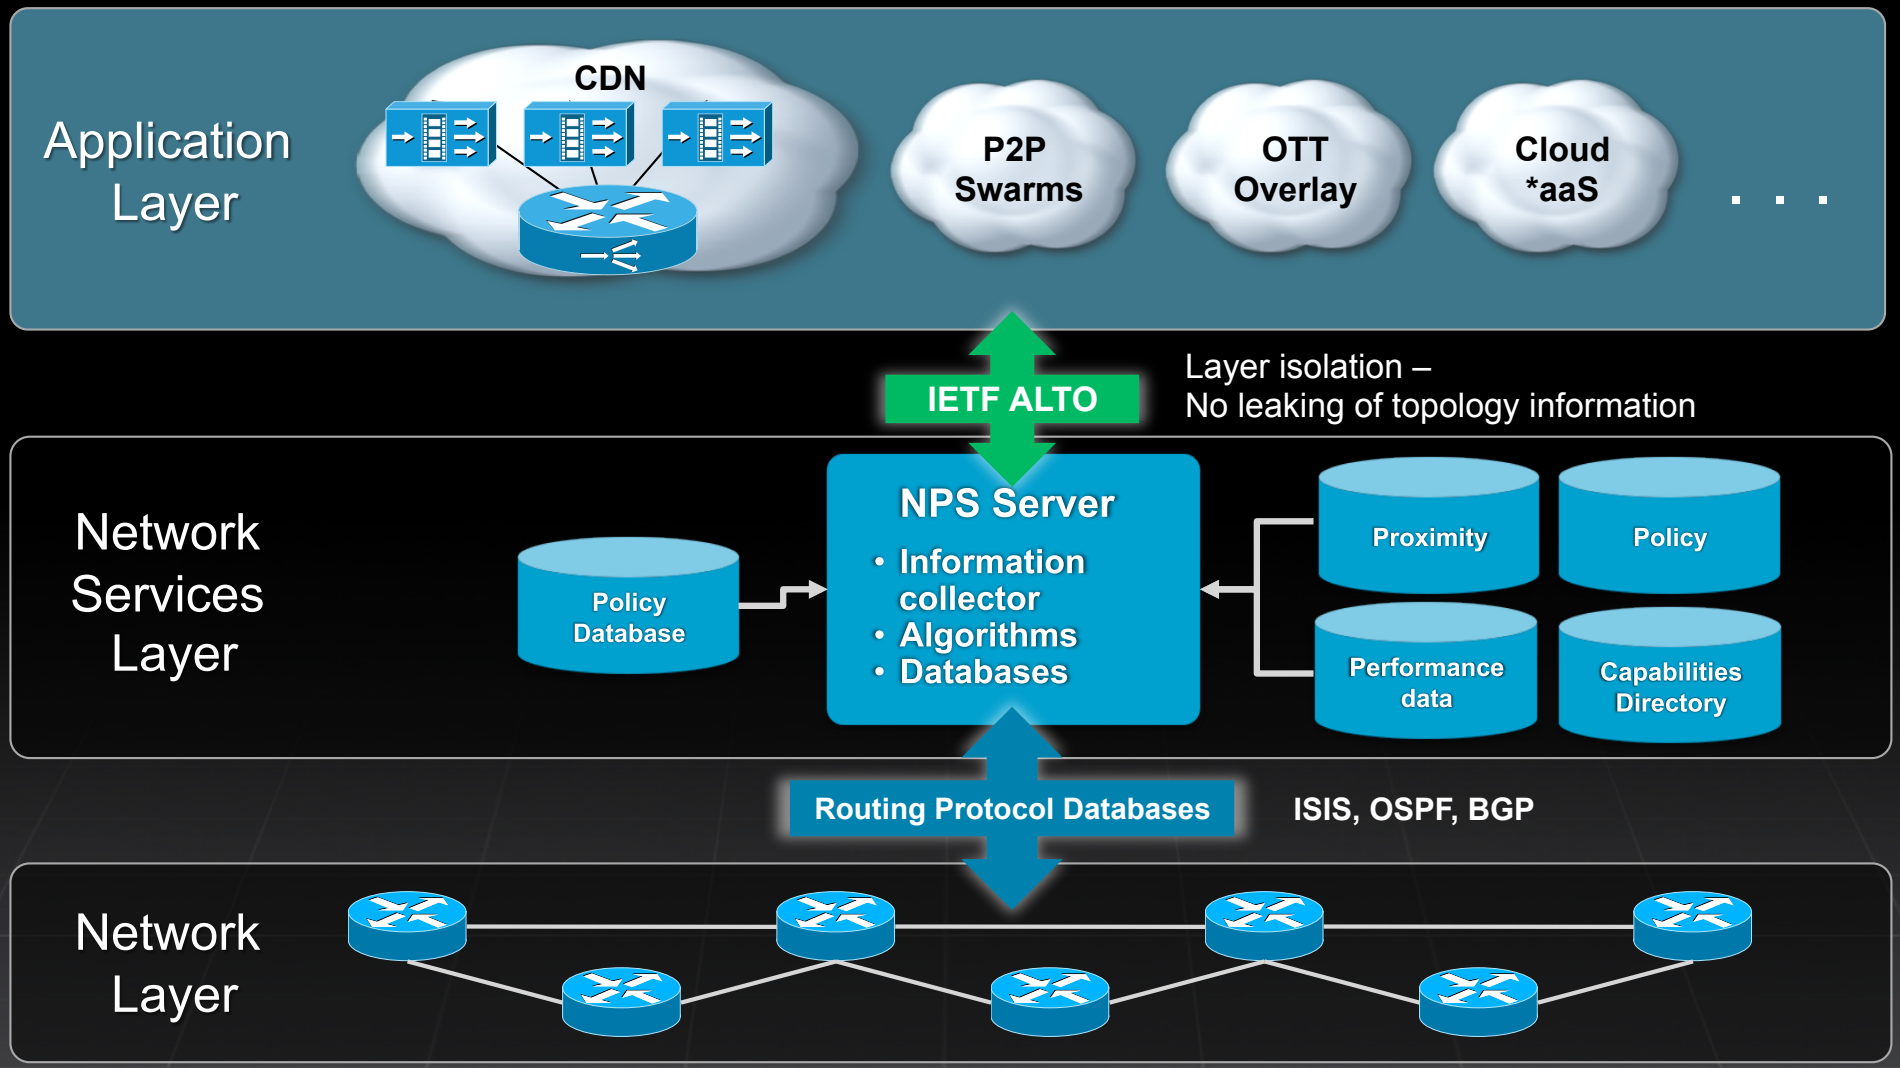

Use Case: Content Delivery Network

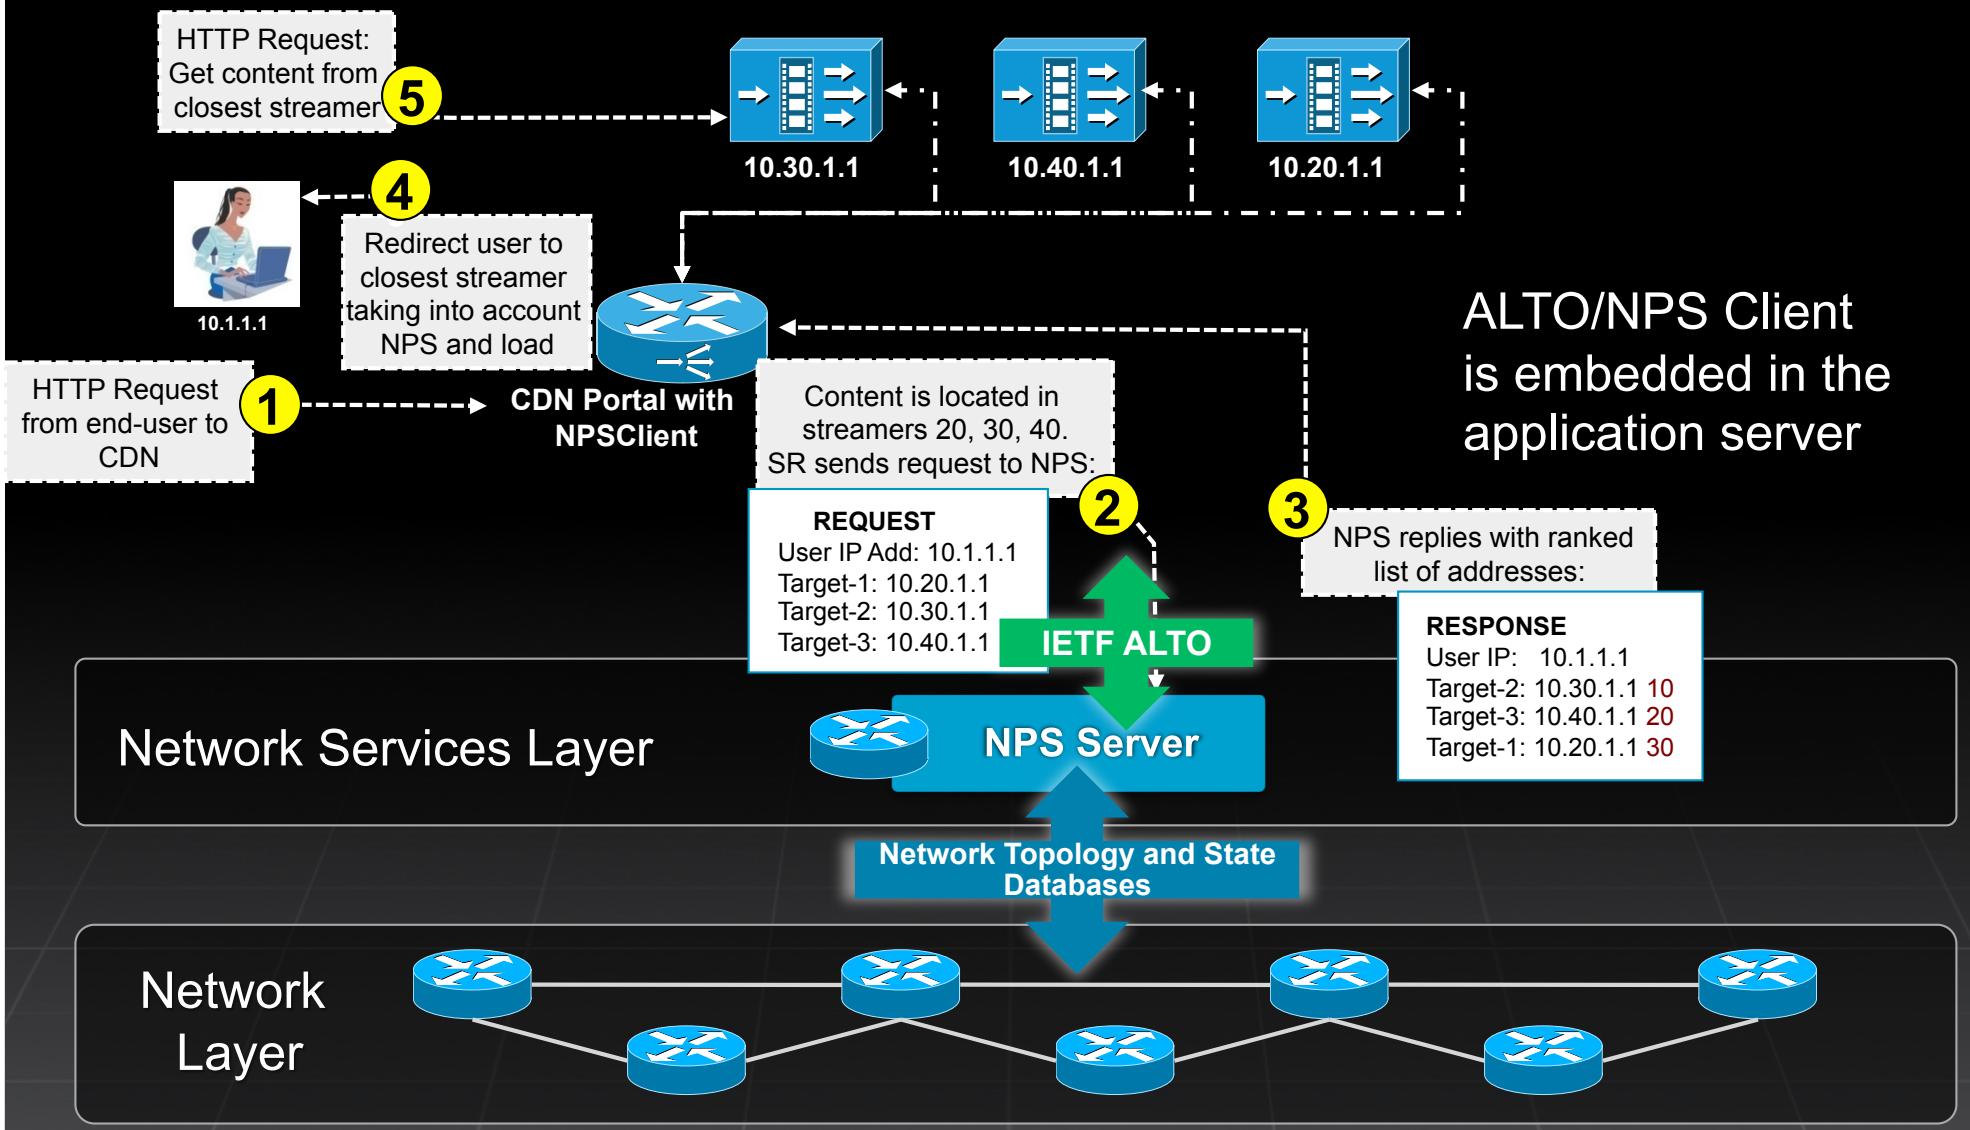

ALTO/NPS Client is embedded in the application server

DNS Resolution Enhanced with NPS

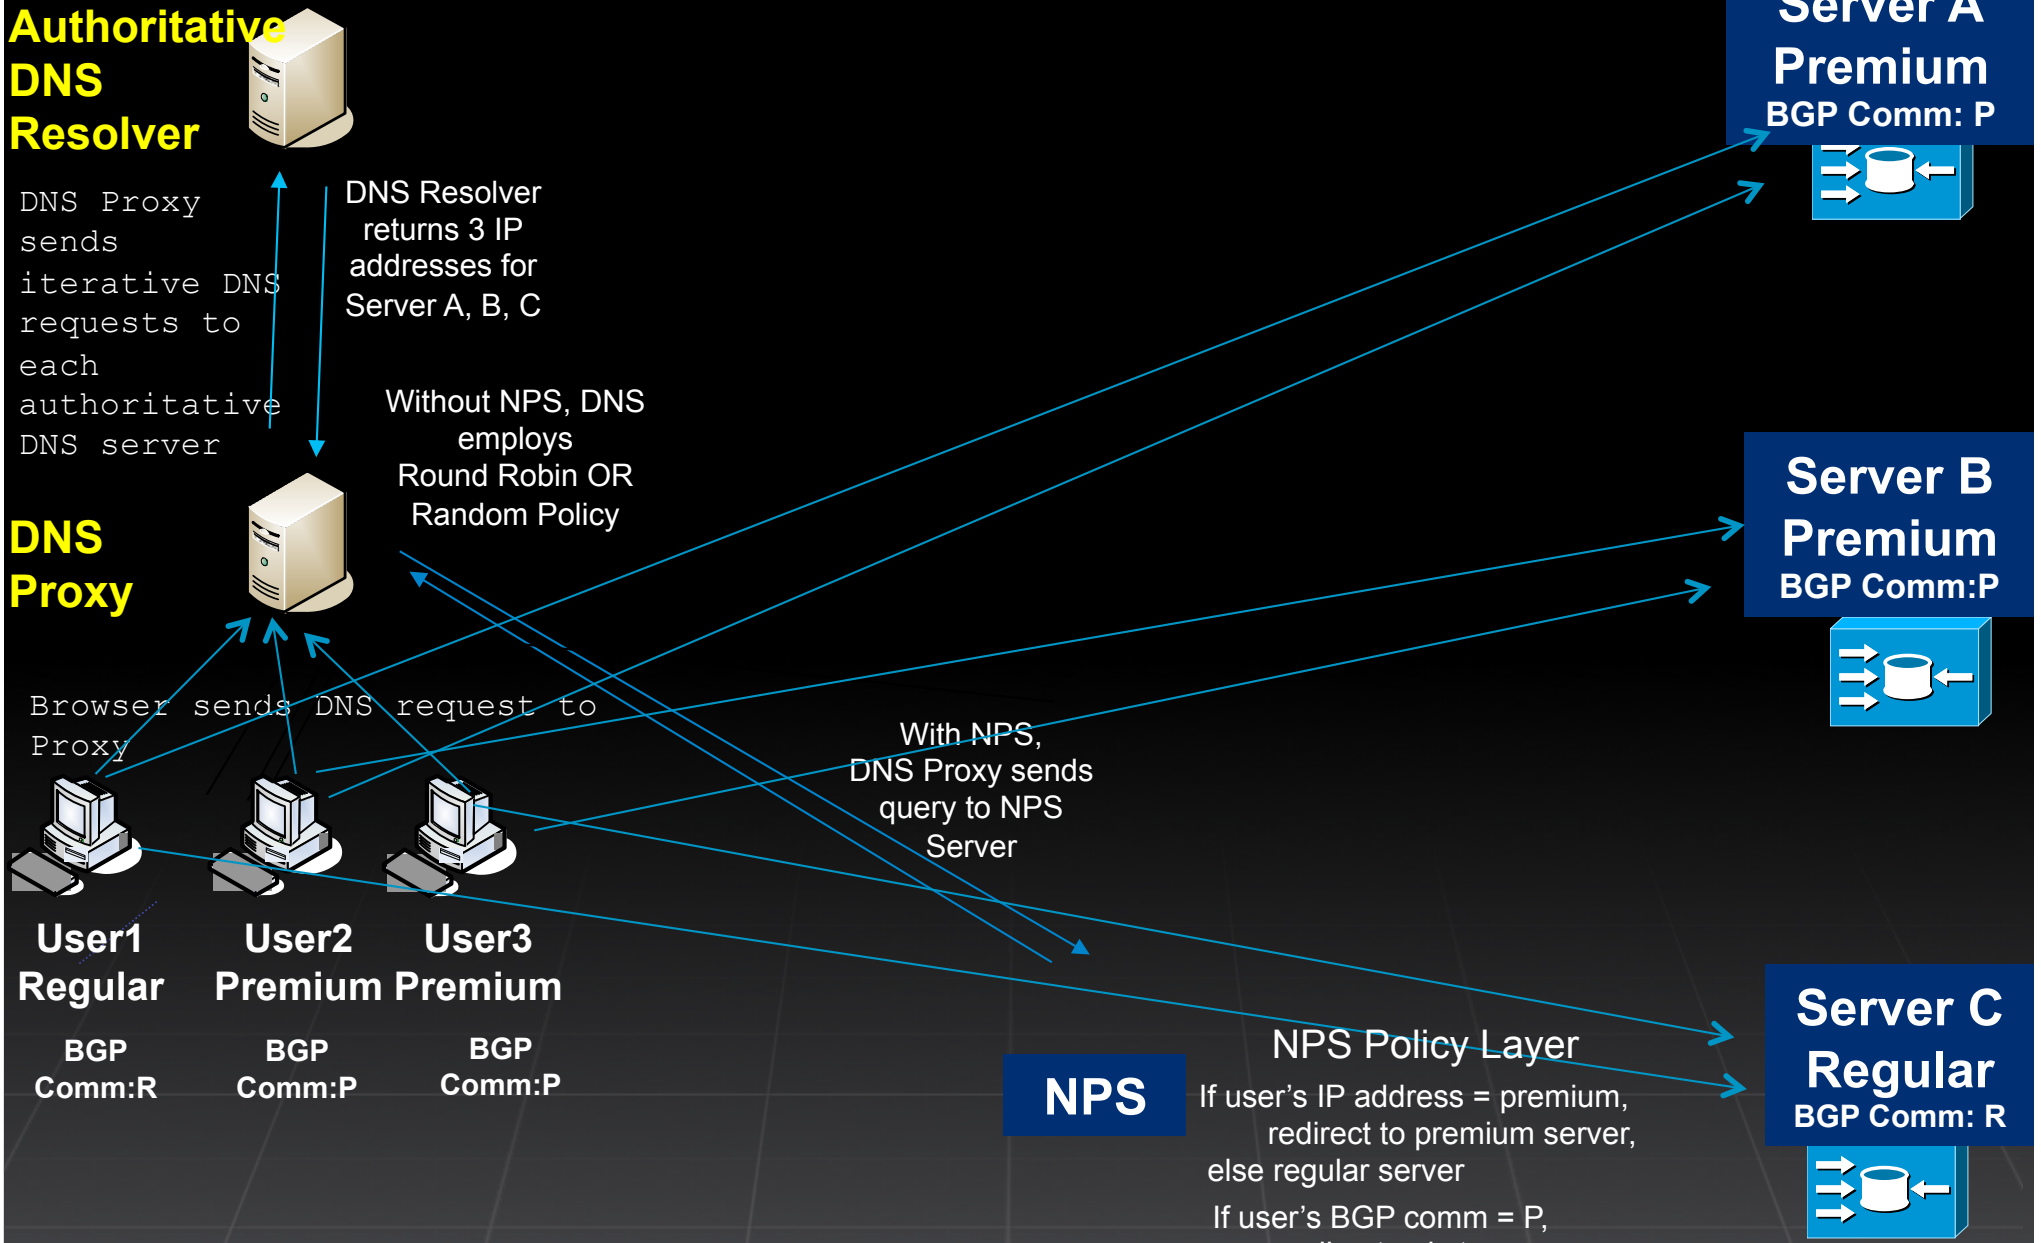

Client enters URL into browser
http://www.video.com
DNS Proxy returns A record IP address to browser

NPS : Benefits of Integrated-in-Router

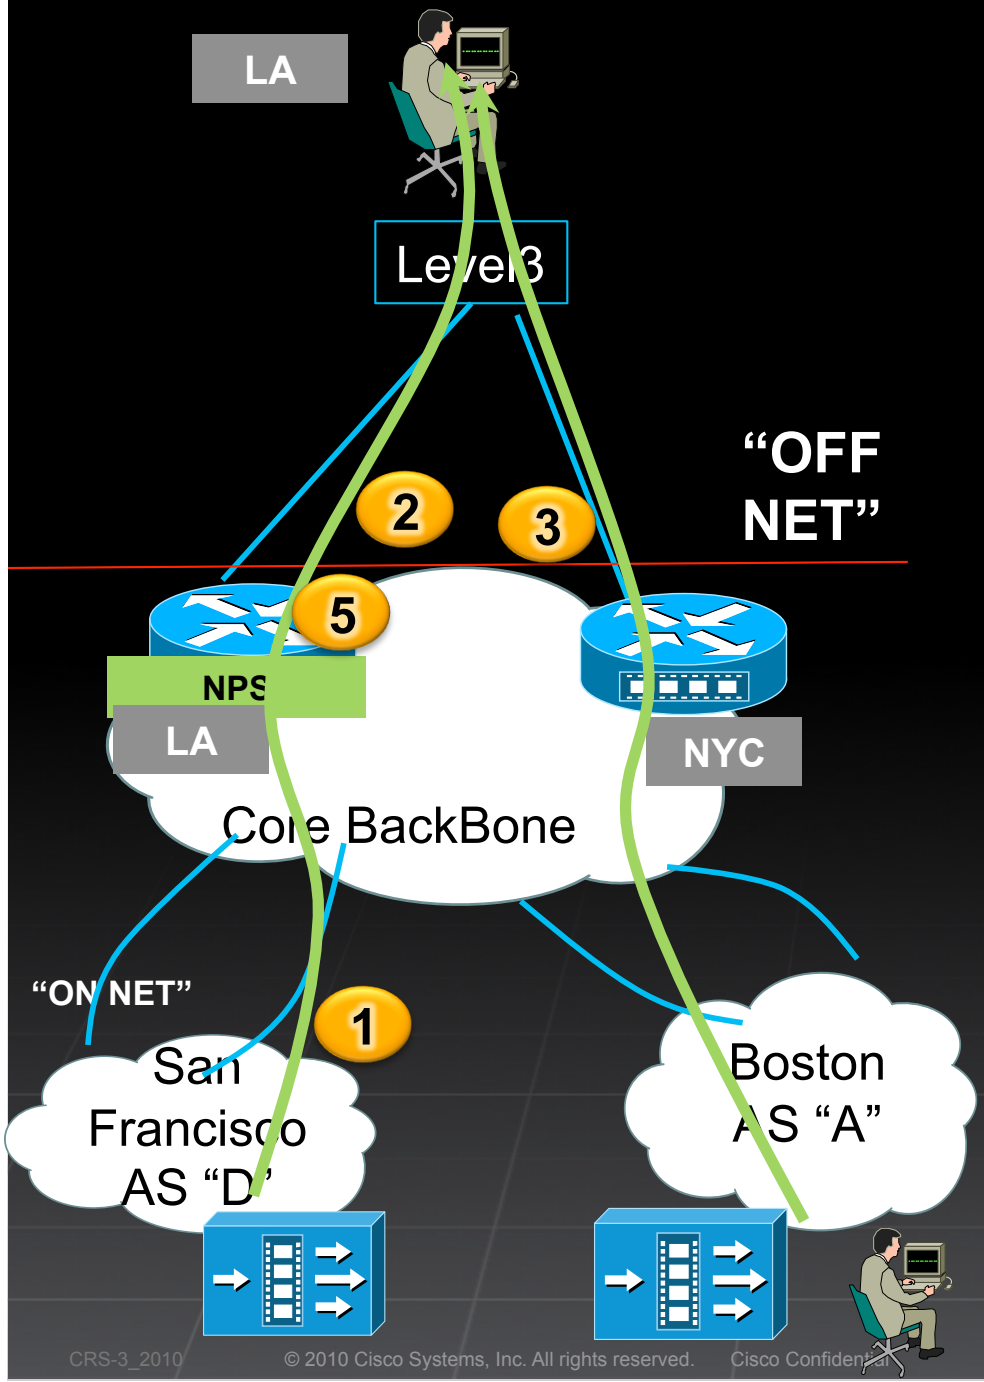

1. LA User served content from LA streamer cache
2. Bandwidth used exceeds negotiated link bandwidth SLA
3. NPS utilizes link bandwidth to change proximity decision, Boston streamer cache now preferred
4. Bandwidth Overage Expenses avoided
5. Time-of-day policy to influence proximity decision
   During day prefer west link,
   During night prefer east link

Use Case: Peer-to-peer overlays

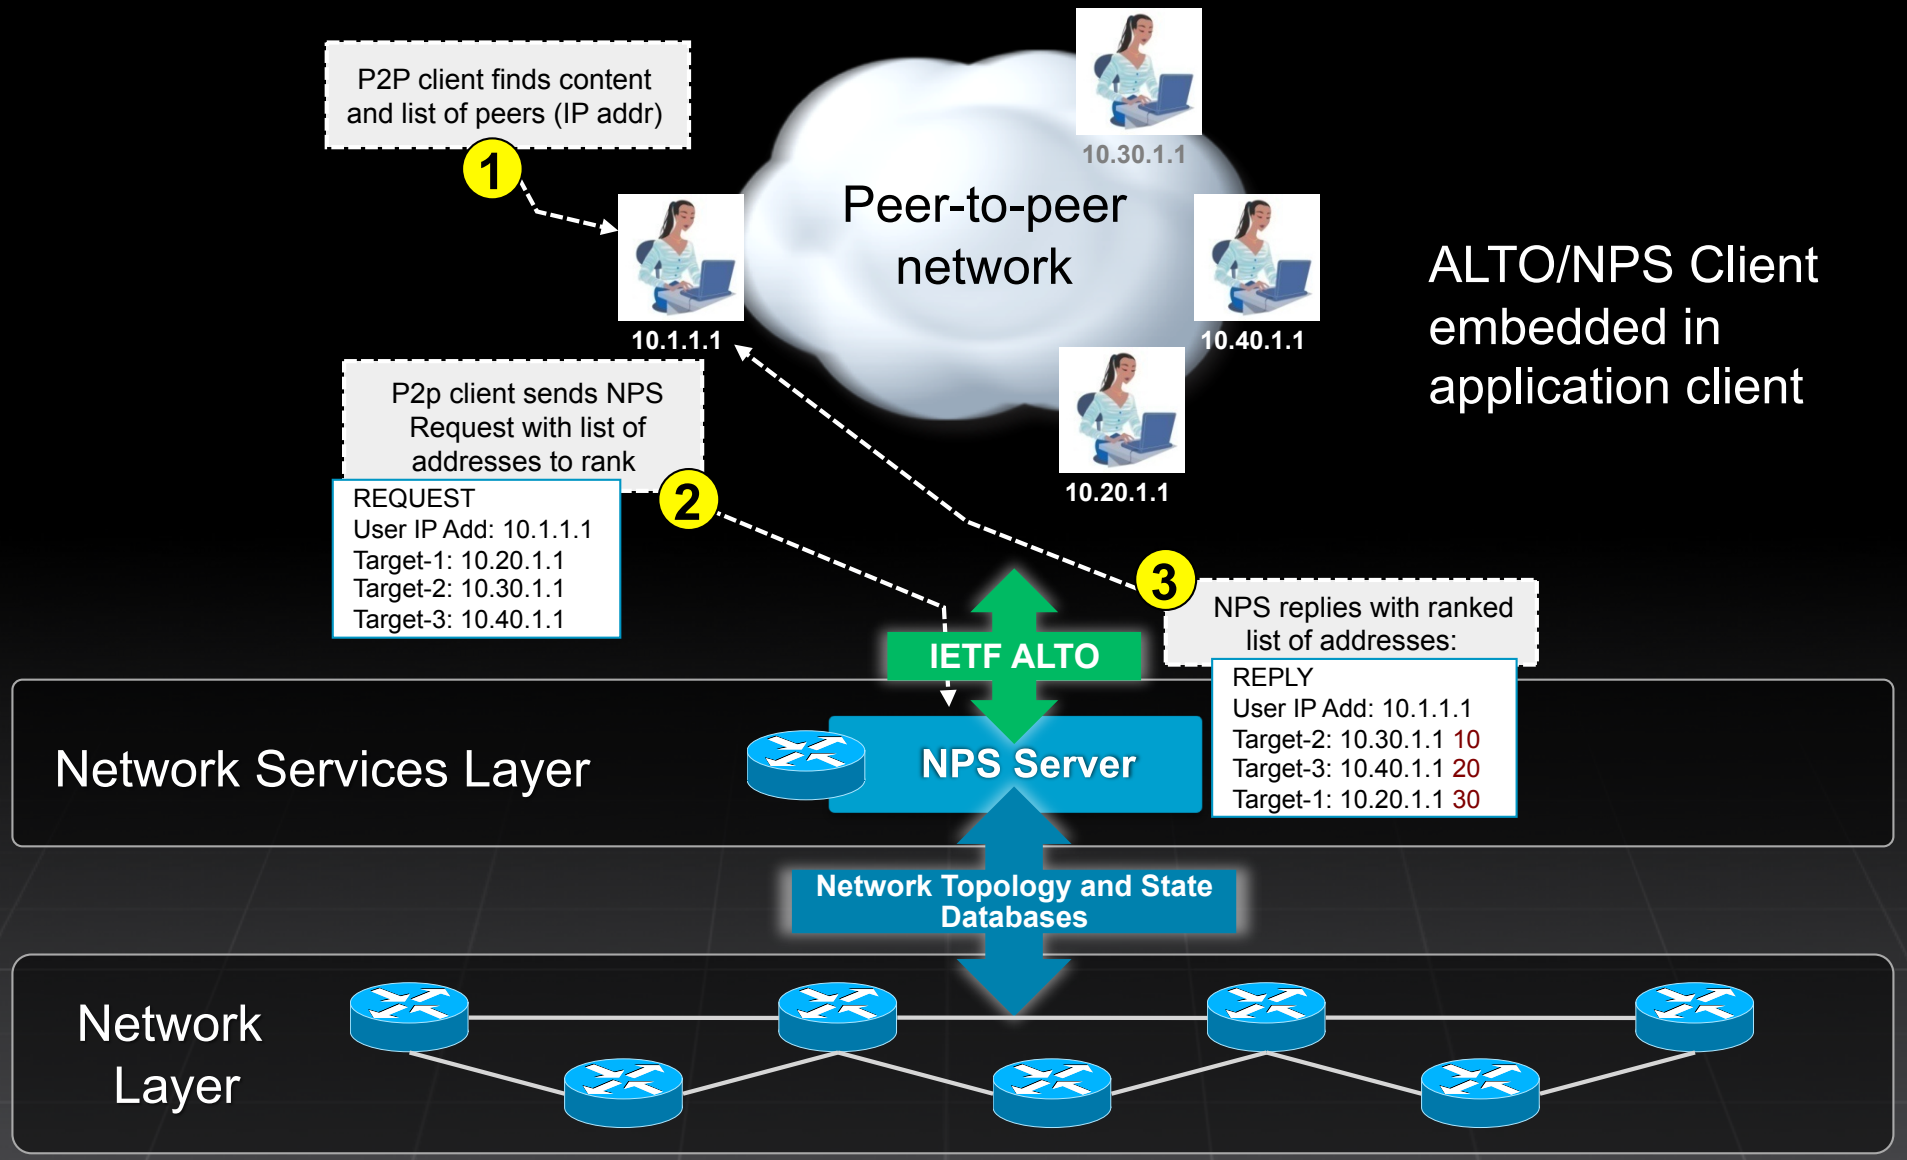

ALTO/NPS Client embedded in application client

Benefits of closer coupling

- When the overlay topology is network aware, it is highly correlated with the underlying network topology;the nodes within an AS form a dense cluster, with only a few connections going to nodes in other AS [3]
- Comcast's experience:
  - *"... reduced outgoing Internet traffic by an average of 34% at peering points."*
  - *"... reduced incoming Internet traffic by an average of 80% at peering points."*

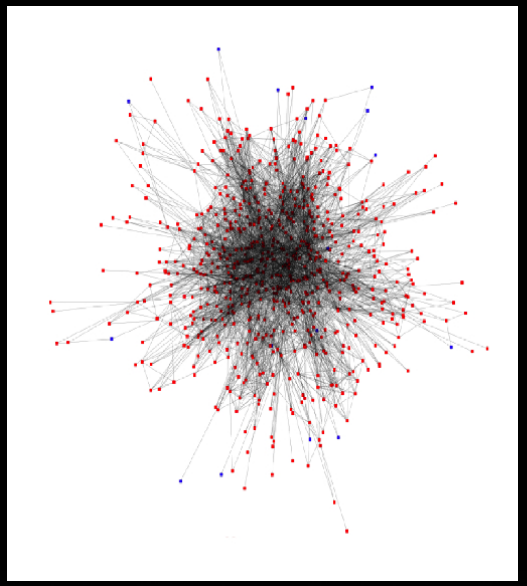

Network Unaware

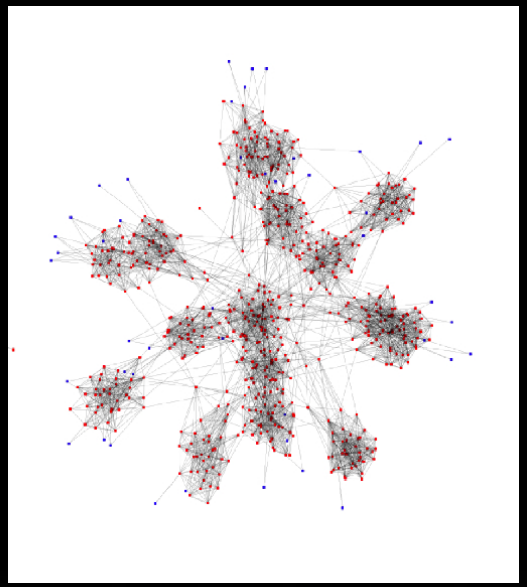

Network Aware (Overlay-underlay Topology Correlation)

[3] Aggarwal, V., Feldmann, A., and C. Scheideler, "Can ISPs and P2P systems co-operate for improved performance?", ACMSIGCOMM Computer Communications Review (CCR), 37:3, pp. 29-40.

[4] C. Griffiths, J. Livingood, L. Popkin, R. Woundy, Y. Yang, "Comcast's ISP Experiences in a Proactive Network Provider Participation for P2P (P4P) Technical Trial", RFC 5632, September 2009

Why is NPS useful even with Akamai/Limelight

1. User requests content
2. Content located in cache
   Request serviced from cache
3. Local POP Cache Overwhelmed OR User requested content NOT located in cache
4. Akamai chooses Upstream Cache independent of SP network policies
   SFO-BOS should carry Premium traffic only
   Don't use expensive transit links for non-revenue-generating-traffic
5. NPS on CRS enables SP to control how Akamai uses SP's network based on SP's traffic patterns

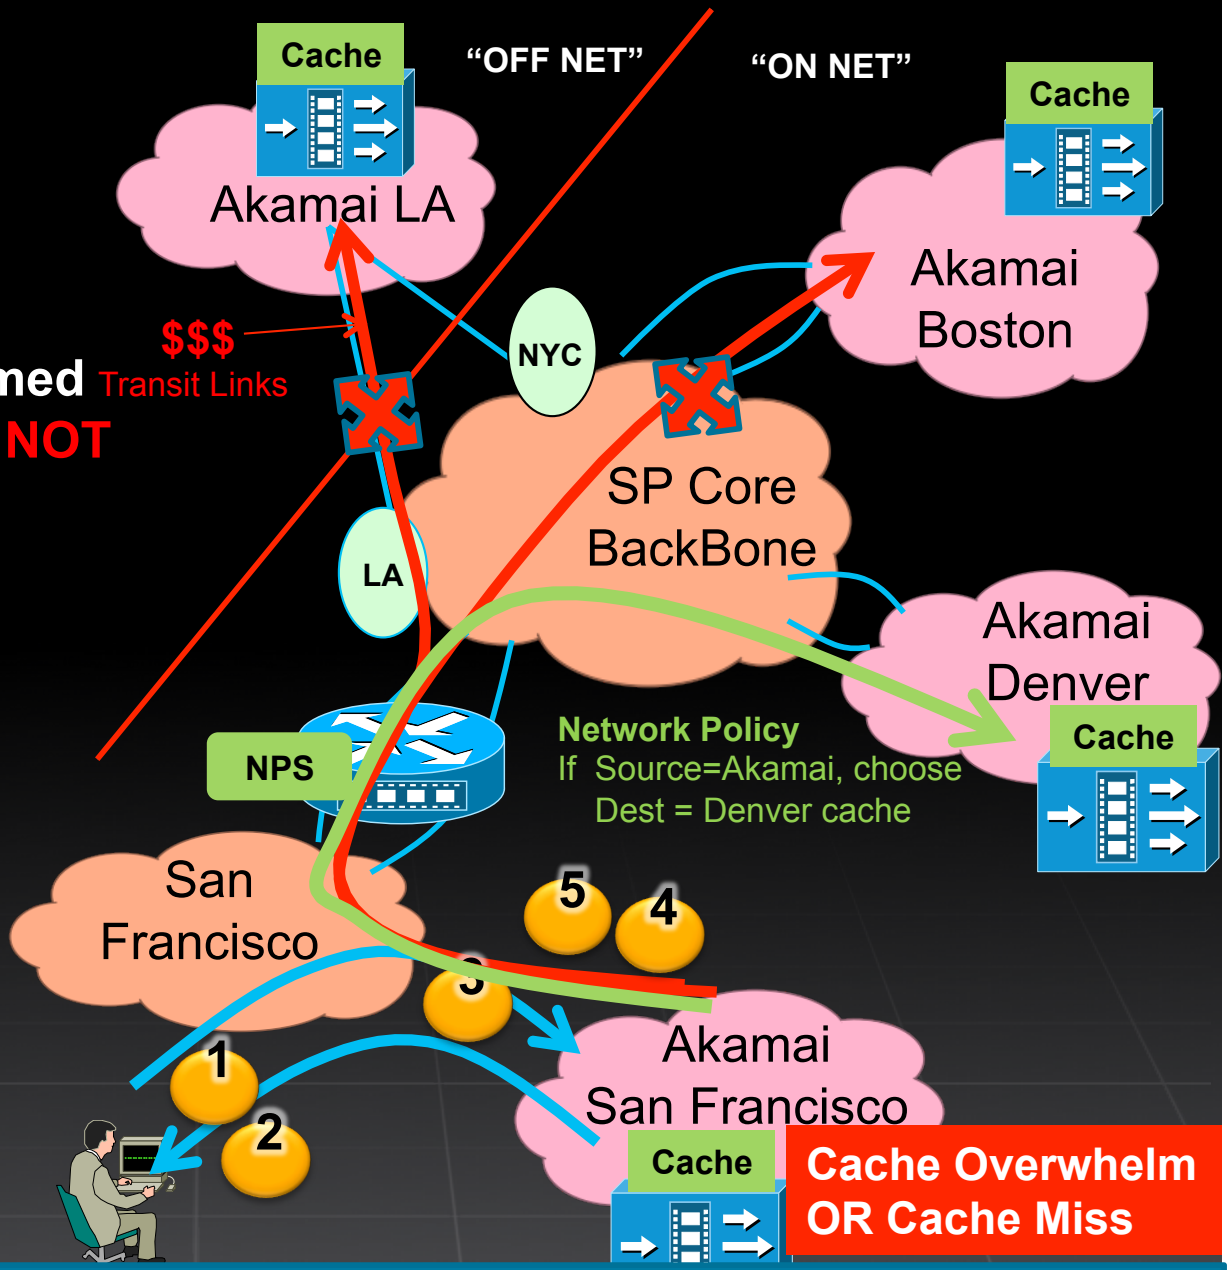

Regain back Control of YOUR Network with NPS

Google Cache Use Case: Enables Premium Service Offering

1. User visits Google.com
2. User ISP's DNS Proxy redirects user to main Google Data Center

   Main Google data center redirects user to caches within SP network

   But Google cannot differentiate premium vs. non-premium users, only User ISP has that knowledge
3. ISP Policy:
   For premium users redirect to closest Google cache
4. For non-premium users redirect to farthest cache or OFF-NET as quickly as possible
5. NPS enables ISP to provide premium level of service for over the top content

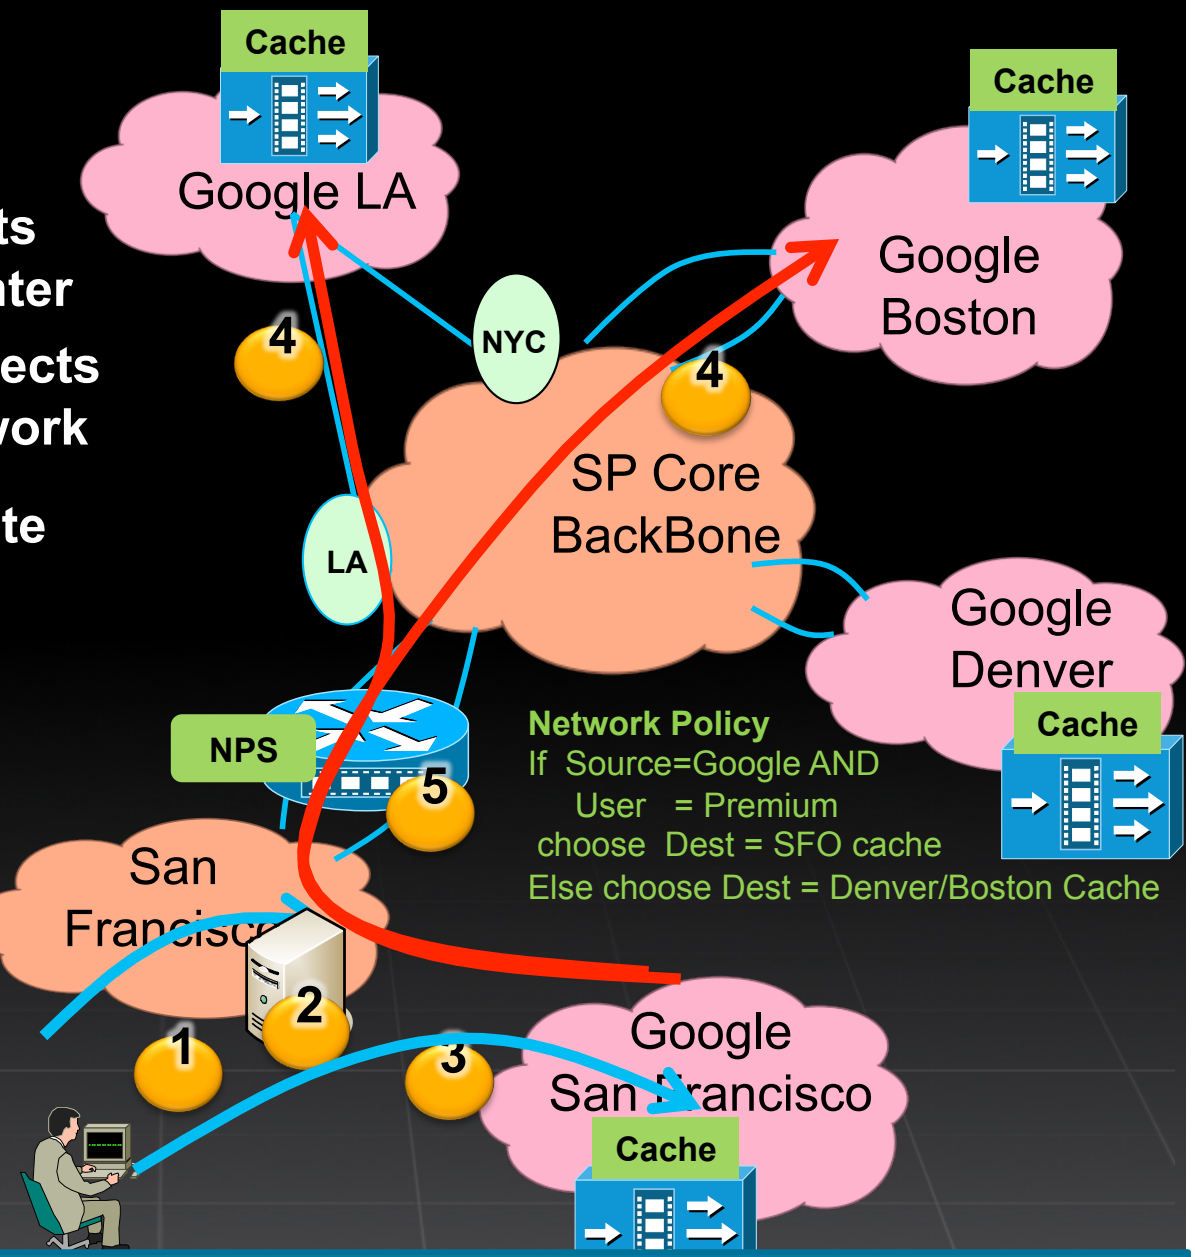

Provide Premium Services for Over the Top Content

Google Cache Use Case #2

1. User visits Google.com
2. User ISP's DNS Proxy redirects user to main Google Data Center
3. Main Google data center redirects user to caches within SP network
4. Google chooses cache regardless of SP network policies

   SFO-BOS should carry Premium traffic only
   Don't use expensive transit links to Google LA for non-revenue-generating-traffic
5. NPS enables SP to control how Google uses SP's network

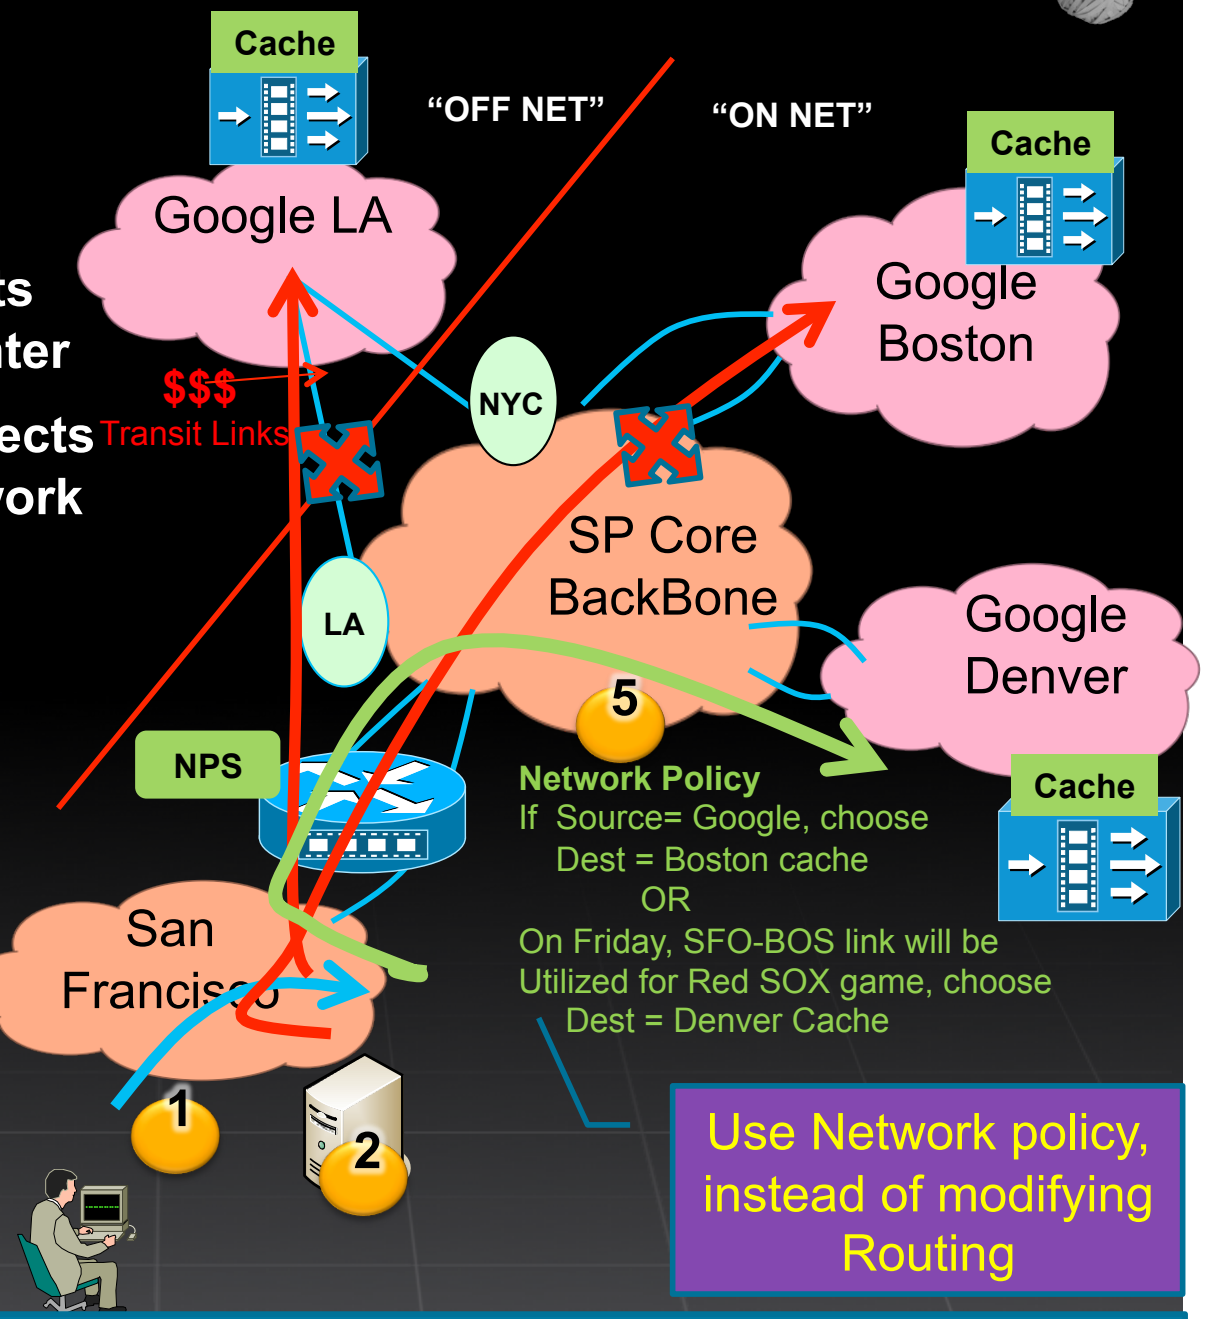

Use Network policy, instead of modifying Routing

Respond to changing traffic patterns with flexibility and agility

Cisco Network Positioning System

- Location definition may override routing visibility
  - Not everything can be grouped through prefix aggregation
- Need for a policy mechanism allowing to group prefixes
  - Good news: it's available and called BGP Communities
- Example: Users in POP1 should first prefer streamers in POP3, then POP2
- Requires:
  - Ability to group prefixes other than through routing paradigm: BGP Community Tagging
  - Ability to define distance/cost/preferences between groups: Policy definition in NPS server

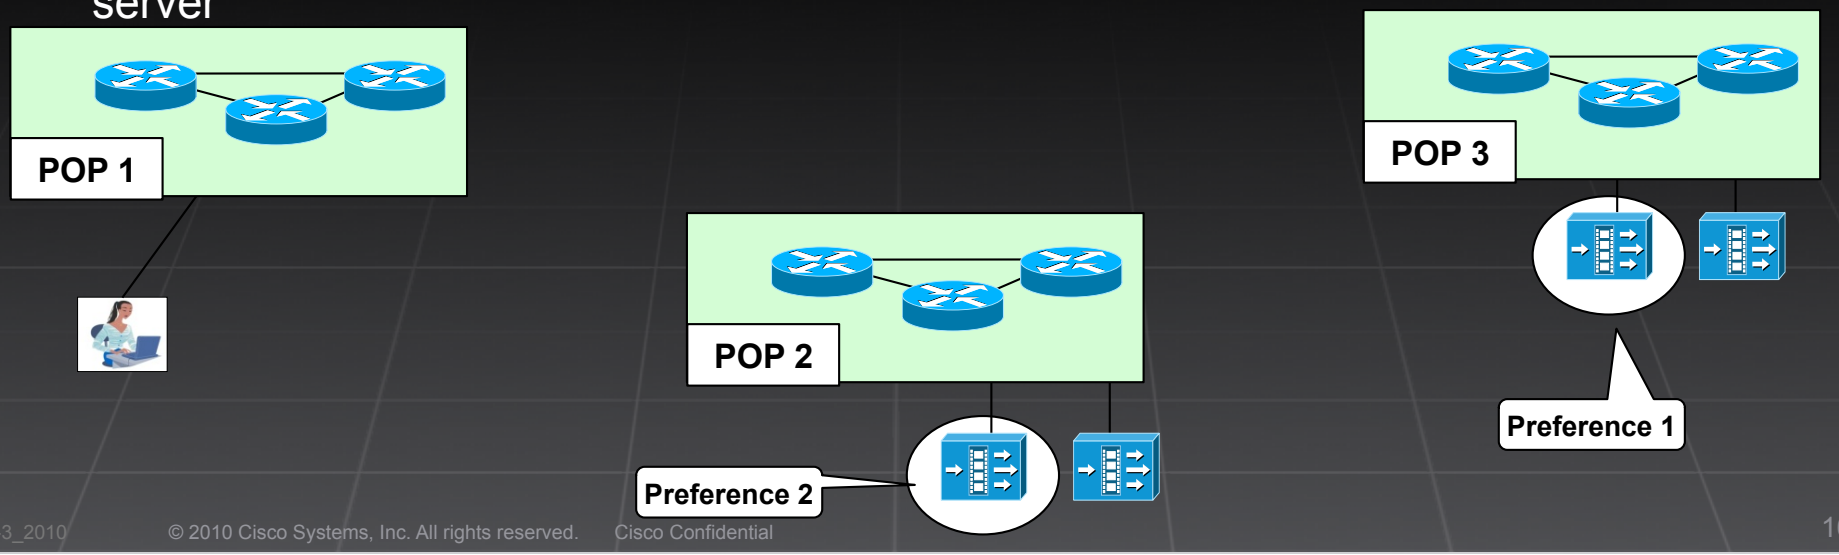

Cisco NPS: Grouping and Policies

- Current implementation:

  NPS co-locates endpoints having prefixes with same BGP Community value

  NPS allow to define arbitrary weight between communities

  Example:

  source-community 111:789 target-community 111:789 weight 5
  source-community 111:789 target-community 111:123 weight 3
  source-community 111:789 target-community 111:456 weight 1

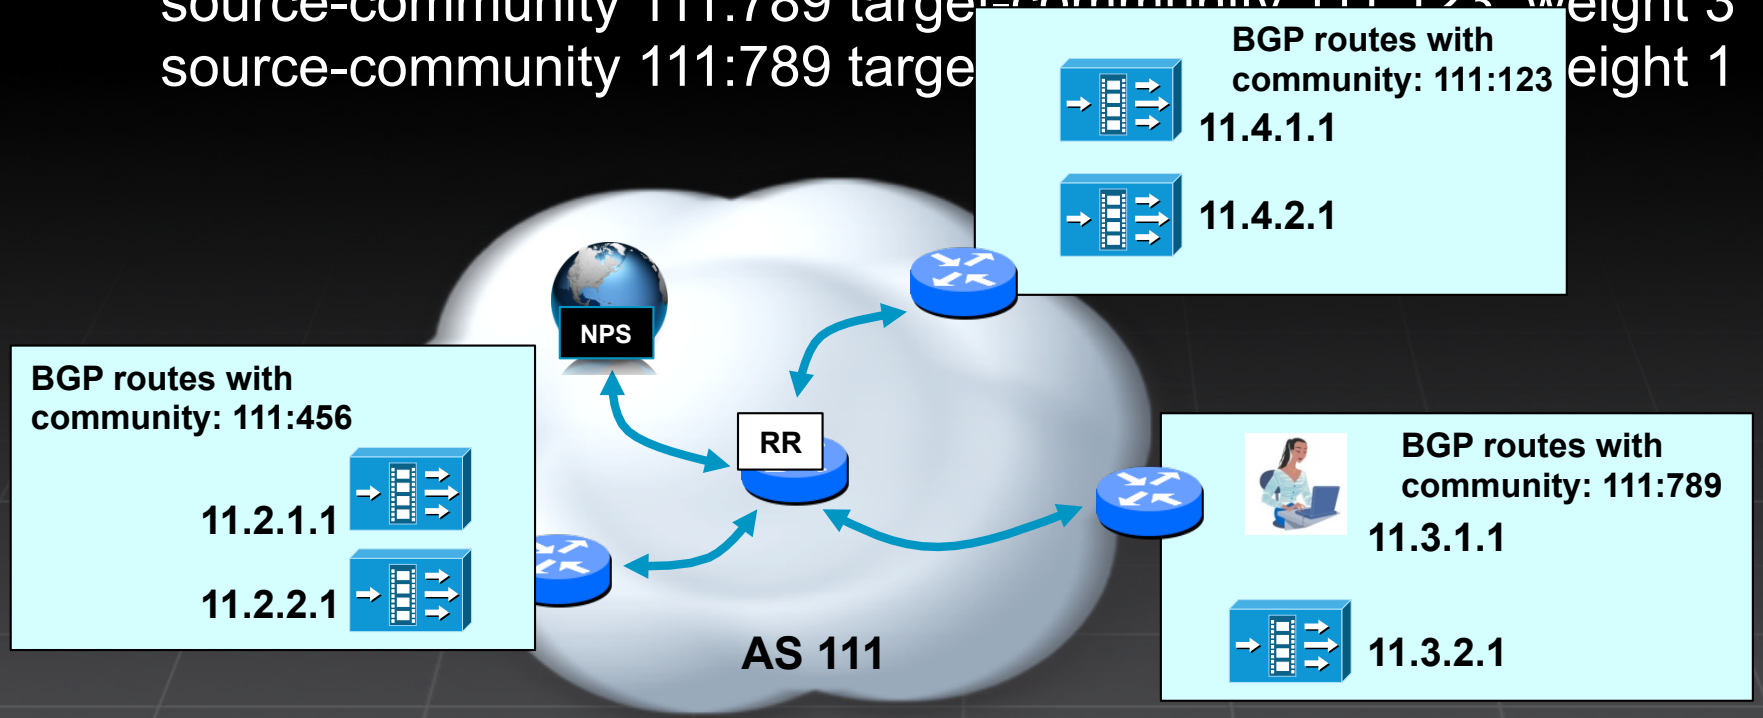

Cisco NPS: Grouping and Policies

- From Topology and Policies to Maps

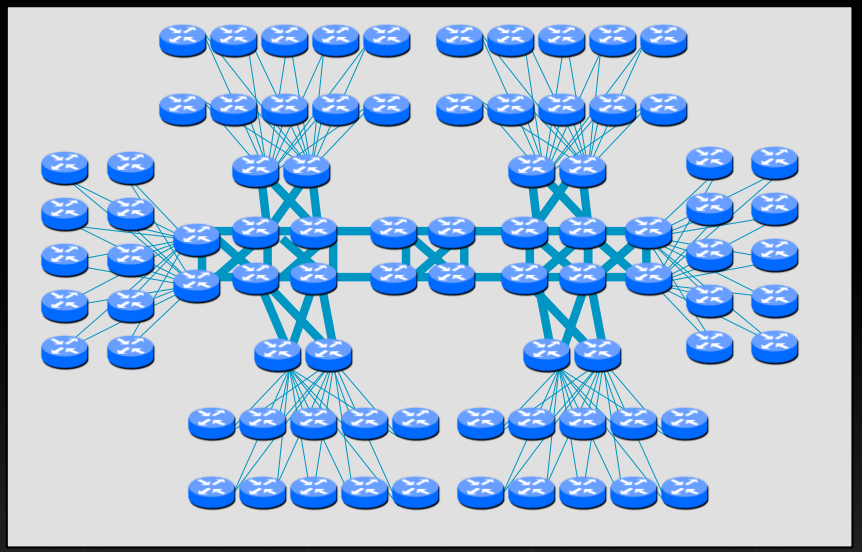

Routing Databases

Policy Databases

State & Performance Data

Location Information

. . .

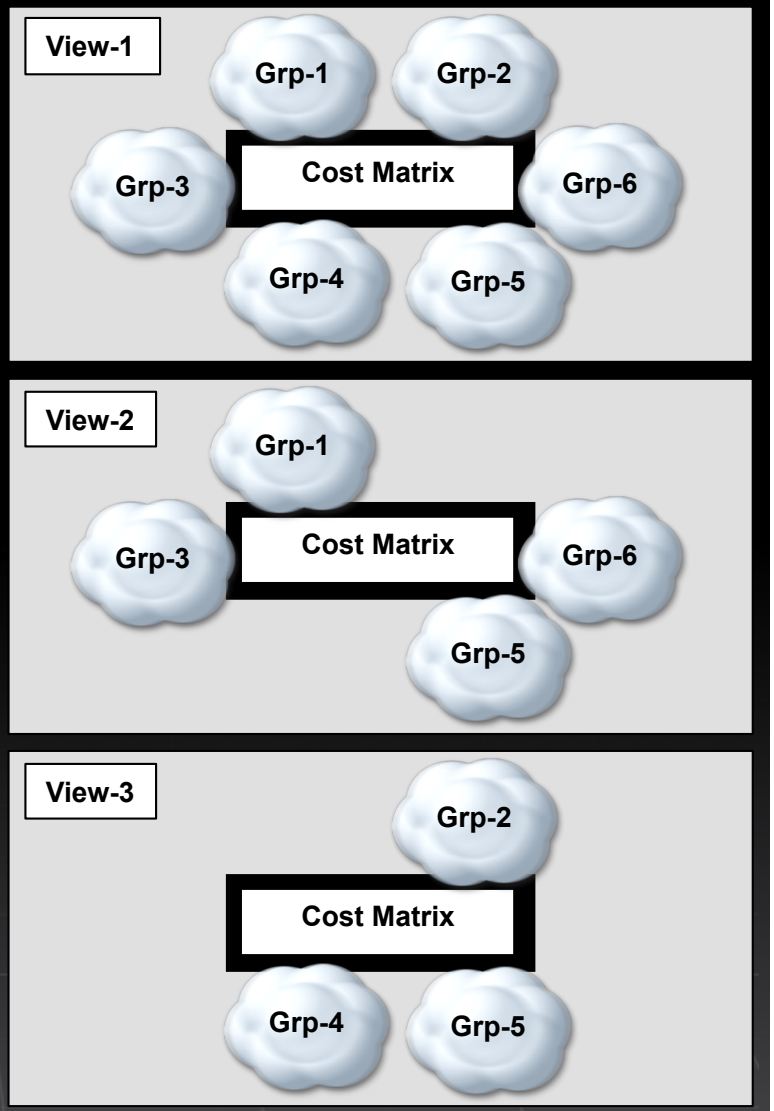

Network Positioning System Example: Redirection

- NPS acquires routing information from within the AS
- Requests received within the AS are locally server
- Requests received for addresses outside the AS will be re-directed to NPS server located in addresses' AS
- NPS servers exchange info through an application overlay

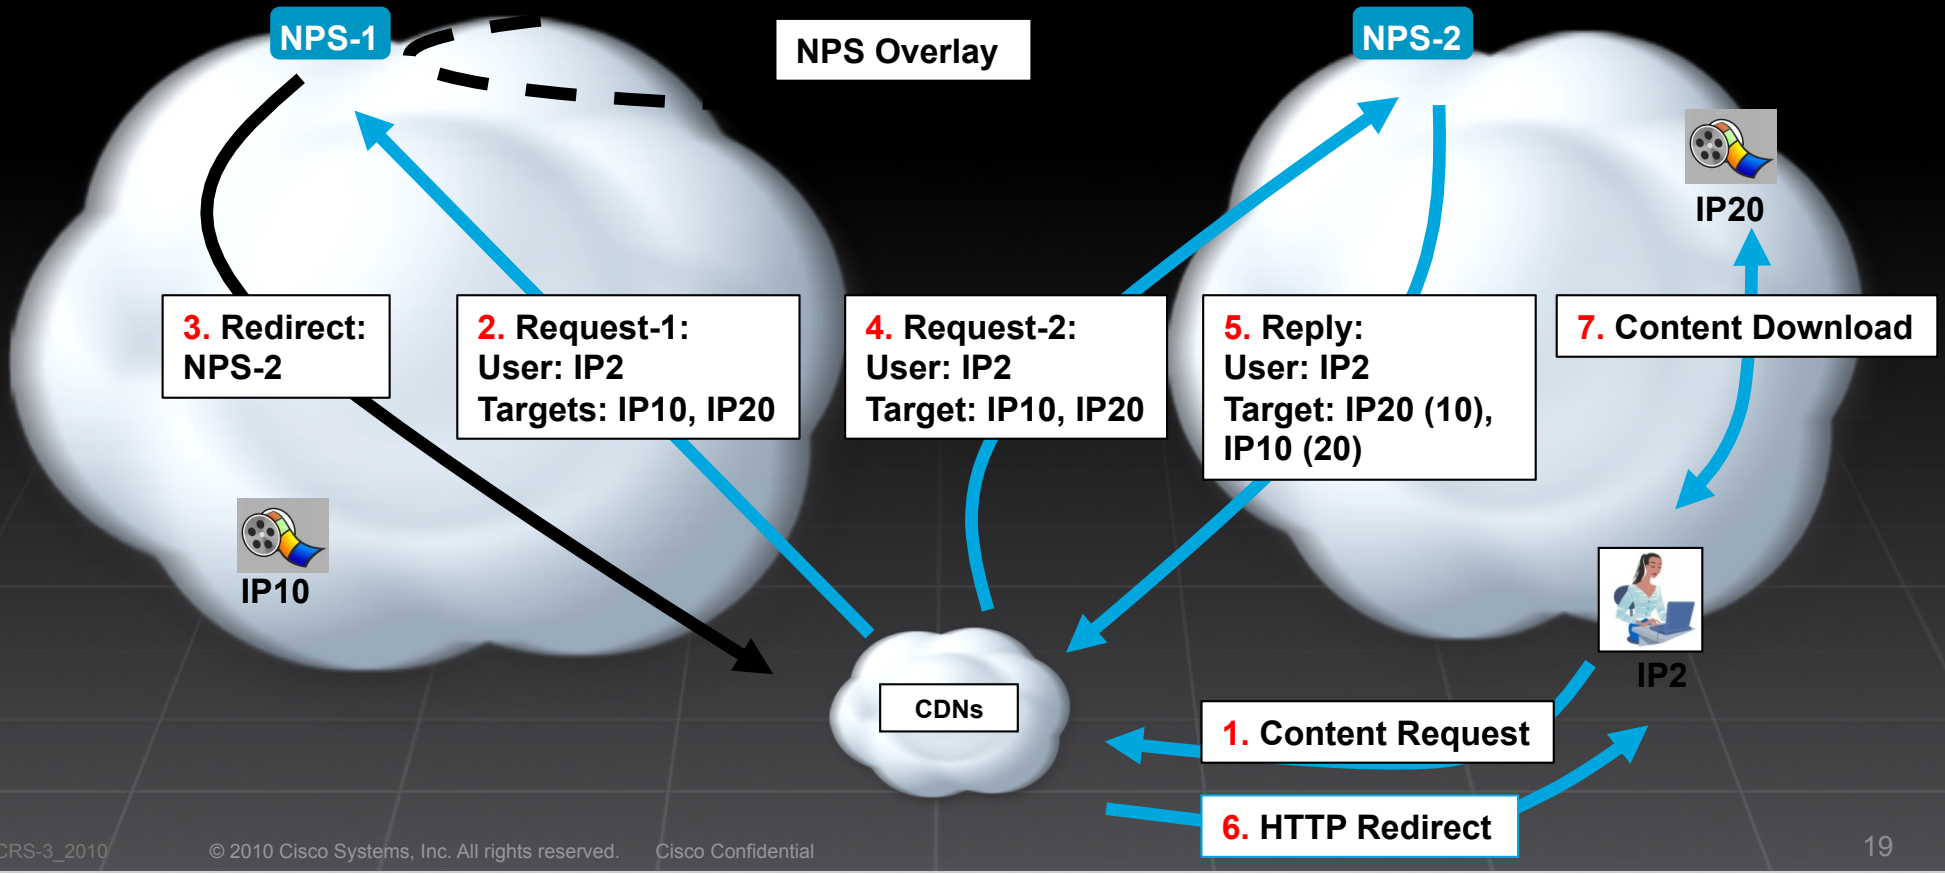

Questions?

Any other use case ?

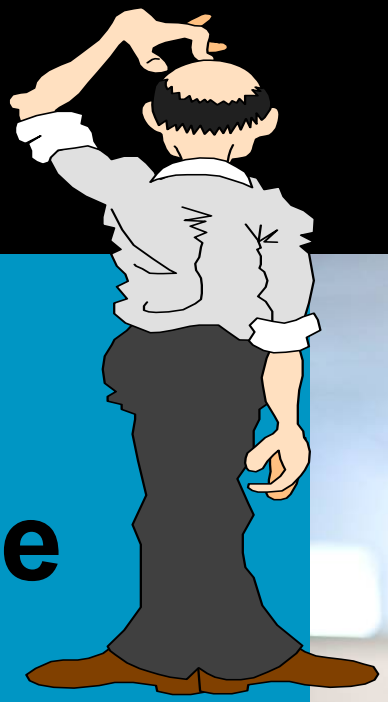

Questions?

Any other use cases?

Backup

Network Positioning System (NPS)

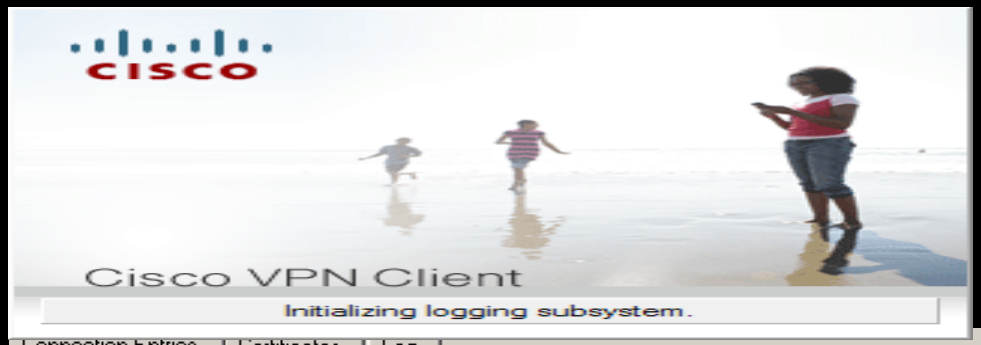

Connection Entries | Certificates | Log

Connection Entry	Host	Transport
Tokyo	tky-vpn-cluster.cisco.com	IPSec/UDP
Tel Aviv	tlv-vpn-cluster.cisco.com	IPSec/UDP
Sydney	syd-vpn-cluster.cisco.com	IPSec/UDP
Singapore	sin-vpn-cluster.cisco.com	IPSec/UDP
Shanghai	shanghai-vpn-cluster.cisco.com	IPSec/UDP
San Jose	**sjc-vpn-cluster.cisco.com**	**IPSec/UDP**
RTP	rtp-vpn-cluster.cisco.com	IPSec/UDP
Richardson	rcdn-vpn-cluster.cisco.com	IPSec/UDP
Johannesburg	jhb-vpn-cluster.cisco.com	IPSec/UDP
Hong Kong-2	hkidc-vpn-cluster.cisco.com	IPSec/TCP
Hong Kong-1	hkidc-vpn-cluster.cisco.com	IPSec
Hong Kong	hkidc-vpn-cluster.cisco.com	IPSec/UDP
GGSG-RTP	rtp-ggsgvpn-cluster.cisco.com	IPSec/UDP
CRDC-Shanghai	shanghai-crdcvpn-cluster.cisco.com	IPSec/UDP
CRDC-San Jose-2	sjc-crdcvpn-cluster.cisco.com	IPSec/TCP
CRDC-San Jose-1	sjc-crdcvpn-cluster.cisco.com	IPSec
CRDC-San Jose	sjc-crdcvpn-cluster.cisco.com	IPSec/UDP
CRDC-Hong Kong-2	hkidc-crdcvpn-cluster.cisco.com	IPSec/TCP
CRDC-Hong Kong-1	hkidc-crdcvpn-cluster.cisco.com	IPSec
CRDC-Hong Kong	hkidc-crdcvpn-cluster.cisco.com	IPSec/UDP
Boxborough	bxb-vpn-cluster.cisco.com	IPSec/UDP
Bangalore	bangalore-vpn-cluster.cisco.com	IPSec/UDP
Amsterdam	ams-vpn-cluster.cisco.com	IPSec/UDP

What's wrong with this picture?

Geographic Closest = best?

Network Closeness =best

Who has Control:

SP or User?

Who should have Control:

SP or User?

Who is Impacted by the decision: SP, User?

User: Experience

SP : $$$$$$

Who has the Intelligence:

Network or Application?

Who is deriving most value:

Network or Application?