Using the Service Control Engine and Deep Packet Inspection in the Data Center

This document is a basic primer for use and deployment of the Service Control Engine products in the data center. This document also provides a basis for further discussion of deep packet inspection (DPI) techniques and products.

Contents

Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Introduction

Deep packet inspection (DPI) provides the ability to look into the packet past the basic header information. DPI intelligently determines the contents of a particular packet, and then either records that information for statistical purposes or performs an action on the packet.

Applications enabled by DPI include the following:

- Traffic Management, or the ability to control end-user applications such as peer-to-peer applications
- Security, resource, and admission control
- Policy enforcement and service enhancements such as personalization of content or content filtering

Benefits include increased visibility into the network traffic, which enables network operators to understand usage patterns and to correlate network performance information along with providing usage base billing or even acceptable usage monitoring.

DPI can also reduce the overall costs on the network by reducing operation expenses (OpEx) and capital expenses (CapEx) by providing a more thorough understanding of what is happening with the network, and by providing the ability to direct traffic or to prioritize traffic more intelligently.

Cisco currently has two hardware-based solutions for achieving this DPI functionality: the Cisco Service Control Engine (SCE) product line, and the newly-introduced PISA hardware for the Cisco 6500/7600 Supervisor 32. This document provides basic configuration and performance information with regard to the SCE product family as well as providing comparisons between the SCE and PISA products.

Service Control Solution Overview

The Service Control solution requires the following components for implementation.

- Management network

  The Service Control solution requires that a management network be created so that the SCEs, Collection Manager, Service Control Application Suite, and the Service Control Application Suite Reporter can communicate with each other. This network can be separate, or it can be part of a larger, more-involved out-of-band management network that many customers deploy to isolate management of network elements from user data.

- Service Control Engine

  The SCE is purpose-built hardware platform that performs deep packet inspection, identifies the users, and generates the report data records. The SCE 2000 provides four Gigabit Ethernet (SX or LX) user and network ports, and two FastEthernet management ports. This device requires that a FastEthernet connection be made to the management network, and the unit must see mirrored/duplicated/couples traffic bi-directionally. This document focuses on the SCE 2000 product. (See Figure 1 and Table 1.)

Figure 1 SCE 2020

Table 1 SCE Specifications

	SCE1010	**SCE2020**
Interfaces	2-GBE (fiber SX/LX)	4-GBE (fiber SX/LX) 4-FE
Management interface	2 x 10/100/1000 FE	2 x 10/100/1000 FE
Processor memory	1.5 GB	1.5 GB
Maximum flows	2M (1M bi-directional)	2M (1M bi-directional)
Maximum subscriber contexts	40,000	80,000
Network configuration	• Receive-only • Inline	• Receive-only • Inline • Cascade

- Collection Manager

  The SCE generates usage data and statistics from the SCE platform and forwards them as raw data records (RDRs), using a simple TCP-based protocol (RDR Protocol). The Service Control solution provides the Collection Manager software as an implementation of a collection system, listening in on RDRs from one or more SCE platforms, and processing them on the local machine. The data is then stored for analysis and reporting functions, as well as simple collection and presentation of data to other operation support and systems (OSS) such as billing. The Collection Manager runs on Sun Solaris or Red Hat Linux platforms. The Collection Manager needs to be connected to the same management network to which the SCE FastEthernet port(s) are connected.

  Table 2 lists the Collection Manager software and hardware requirements.

Table 2 Collection Manager Software and Hardware Requirements

Solaris	**Red Hat Linux**
Hardware	**Hardware**

Table 2 Collection Manager Software and Hardware Requirements (continued)

• Minimum 500 MHz CPU • Minimum 1 GB RAM per CPU • Hard disk: – One hard disk, at least 18 GB (recommended for bundled installations) – A second hard disk (at least 18 GB), to store Sybase data • 100BASE-T network interface	• Minimum 800 MHz CPU • Minimum 1 GB RAM per CPU • Hard disk: – One hard disk, at least 18 GB (recommended for bundled installations). – A second hard disk (at least 18 GB), to store Sybase data • 100BASE-T network interface
Software and Environment	**Software and Environment**
• Solaris 5.8 64-bit build 04/01 or later (currently only Solaris 5.8 and 5.9 is supported). • Solaris core installation The following additional system packages should be installed: • SUNWbash—GNU Bourne-Again shell (bash) • SUNWgzip—GNU Zip (gzip) compression utility • SUNWzip—Info-Zip (zip) compression utility • SUNWlibC—Sun Workshop Compilers Bundled libC • SUNWlibCx—Sun Workshop Bundled 64-bit libC	• Red Hat Linux 3.0 or 4.0. • Red Hat Enterprise “Base” Installation (for installations with bundled Sybase) The following additional package should be installed: • *compat-libstdc++* (available on the Red Hat installation CD.) Latest recommended patches from Red Hat should be applied.

- Service Control Application Suite Broadband (BB) Console

  The SCAS BB console is the SCE GUI used to create, modify, and apply the service configuration. The SCAS BB Console lets you define services, packages, protocols, bandwidth control, and other items in the configuration. The SCAS BB Console creates a policy configuration file (.pqb), which can then be saved and/or applied to the SCE device(s). The SCAS BB Console runs on any Windows PC (Windows 95 or greater), and it requires a connection to the management network.

- Service Control Application Suite Reporter

  The SCAS Reporter allows you to query the Collection Manager RDR database and to present the results in a chart or table. This tool provides a valuable resource for understanding the usage patterns and resources used by applications and users that use network resources. This tool can also help your staff to understand the operational impact of various rules and what their impact might be if they are implemented onto the network. This tool runs on any Windows PC (Windows 95 or greater), and it requires a connection to the management network.

- Subscriber Manager

  This optional but recommended component is required to function as middleware software that is used to bridge between the OSS and the SCE platform(s). Subscriber information is stored in the Subscriber Manager database and can then be distributed between multiple devices according to actual subscriber placement. The Subscriber Manager runs on Sun Solaris or Red Hat Linux platforms. *This element requires a Fast Ethernet connection to the management network.* Software requirements are the same as for the Collection Manager.

The Subscriber Manager provides subscriber awareness, mapping network IDs to subscriber IDs. It obtains subscriber information using dedicated integration modules, which integrate with AAA devices such as RADIUS or DHCP servers.

Subscriber information may be obtained in one of two ways:

- Push mode—The Subscriber Manager pushes subscriber information to the SCE platform automatically upon logon of a subscriber.
- Pull mode—On-demand, in response to a query from the SCE Platform to the Subscriber Manager.

Figure 2 shows a high level connectivity diagram of each of the components of the Service Control solution.

Figure 2 Service Control Solution Components

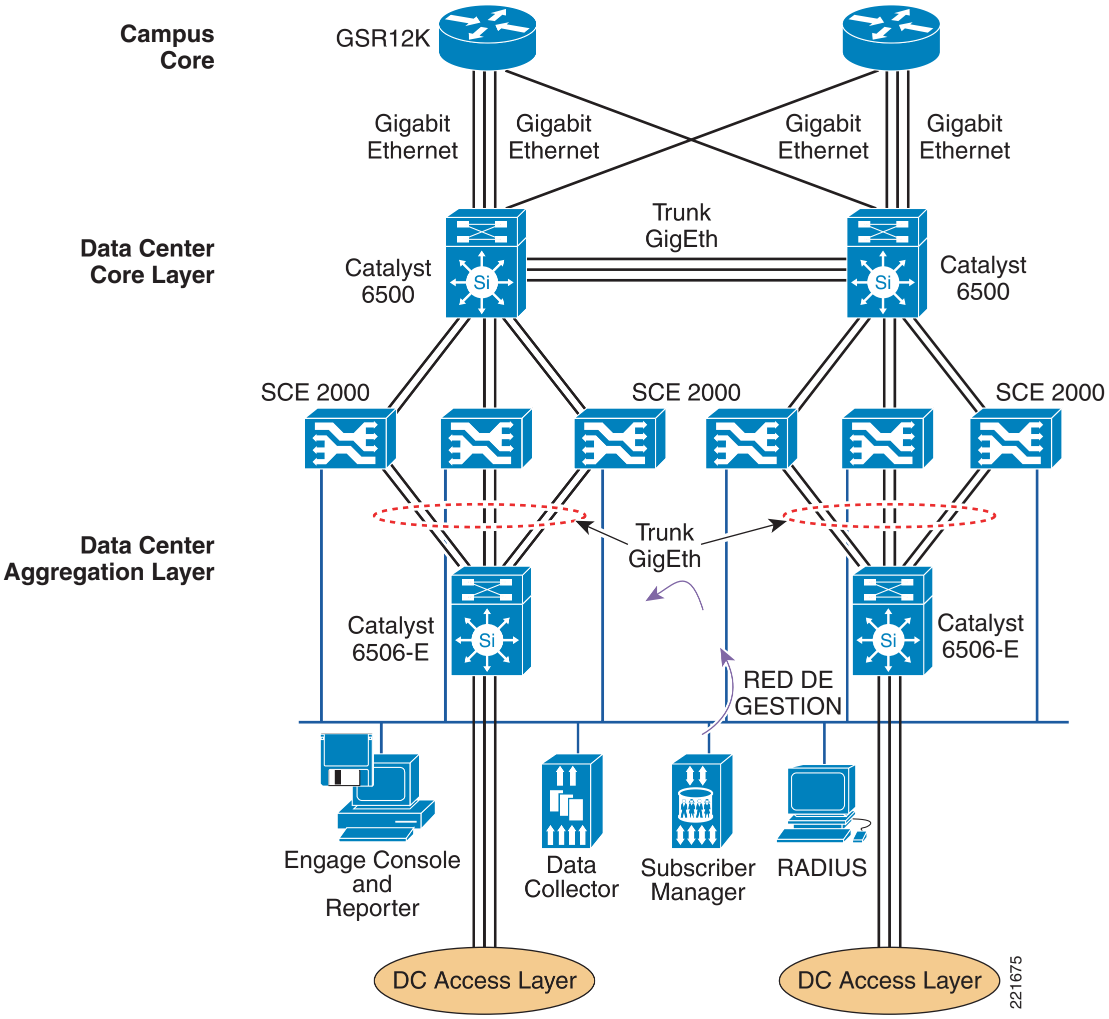

Service Control Engine Insertion Strategies

There are various ways of positioning the SCE products for use in the network. As part of a data center deployment, only a subset of these were used. This is based on the requirements for maintaining Layer 3 functionality between the aggregation and the core layers of the data center network as well as the requirements of moving to 10 Gig throughput rates. The two methods used allow for a customer to first become operationally experienced with the SCE as well as to begin to baseline traffic patterns on their network to begin the policy creation phase. Policy creation should not be attempted without an understanding of the applications on the network and the traffic flow that these applications place on the network.

Note In the planning of your SCE deployment, the SCE needs to see both sides of a traffic flow, and asymmetric traffic flows need to be minimized if proper reporting and policy implementations are to be undertaken. Asymmetric flows may not be reported correctly and can adversely affect performance of created user policies.

Port Mirror

In this deployment scenario, the SCE is deployed in a monitor-only method, and the actual port mirroring can be done at either the core or aggregation layers. This allows the data center operations staff to create the SCE management network and to test functionality, as well as begin to gather baseline traffic analysis on their network.

Inline Multi-Gigabit Service Control Point

In a inline multi-gigabit service control point (MGSCP) configuration, the SCE is inserted between the core data center routers. (See Figure 3.)

Figure 3 Inline MGSCP Deployment

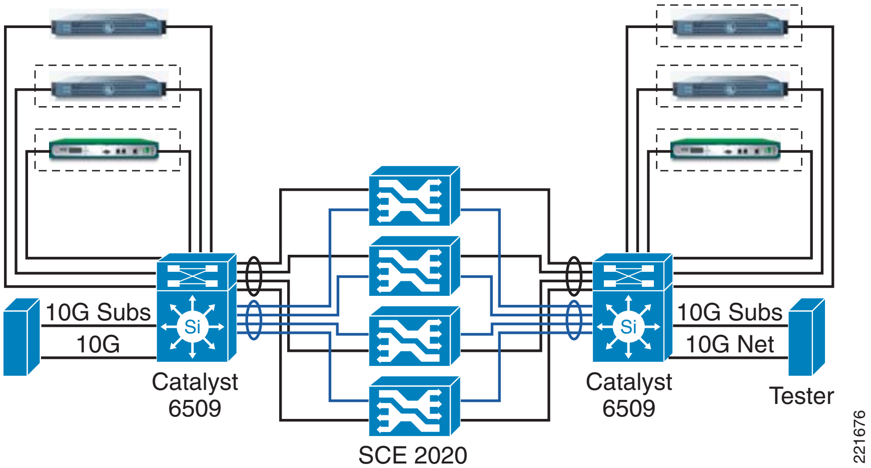

The Service Control solution can also be deployed between the core and aggregation nodes. However, with the requirement of all flows traversing the same SCE, it is more practical to deploy the SCE devices at a layer in the data center and not between. Optional external optical bypass modules and redundant cascade operation provides high availability. The inline deployment permits dynamic control of the traffic including throttling, blocking, and redirection. The benefits of this solution are that any single SCE failure does not result in an outage, and the solution can be architected to provide for an N+1 redundancy model. In addition, from a sizing perspective, this solution allows the data center operators to add additional SCE units as bandwidth needs dictate.

Basics of Dispatch Operation

When configured in dispatch mode, the Catalyst 6500/7600 requires high speed links facing the host device and the network. The Catalyst 6500 dispatches traffic flows to a unique port on the EtherChannel served by an SCE, which performs the configured analysis of the traffic and returns the packets to the original data path. The reason for the unique port is that traffic in the flow must be dispatched to the same SCE to maintain state information.

The Catalyst 6500 gives the network operations staff the flexibility to use Layer 2 or Layer 3 dispatching, depending on their data center deployment. When using Layer 2, packets are dispatched into the EtherChannel using a combination of Layer 2 switching and flow hashing. When Layer 3 routing is used, you need to use policy-based routing (PBR) in combination with flow hashing in the EtherChannel. It is also possible to use both Layer 2 and Layer 3 methods of dispatch to the same SCE cluster group.

For optimum flow dispersion, the number of links in the EtherChannel bundle should be 2^N.

MGSCP Options

Figure 4 shows the Cisco Catalyst 6500/7600 in dispatch mode.

Figure 4 Catalyst 6500/7600 in Dispatch Mode

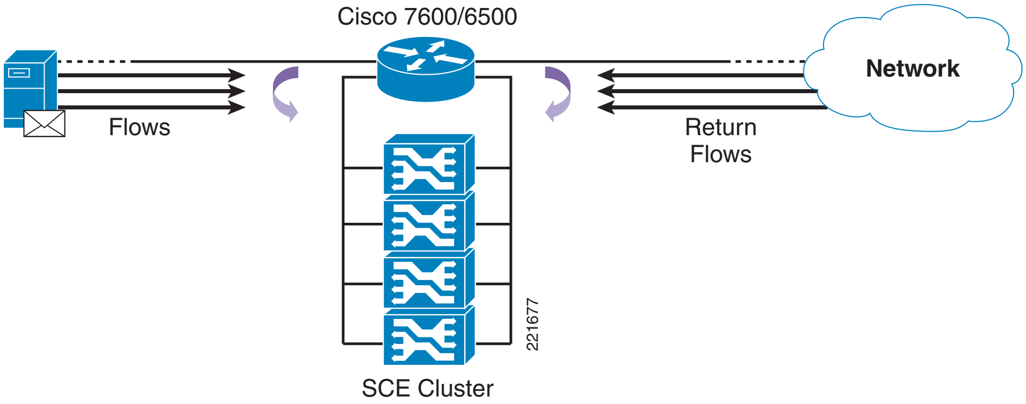

MGSCP Layer 2 Dispatch Mode

When both the host and the network ports are configured as Layer 2 VLAN ports, the MGSCP solution can be deployed in Layer 2 operational mode. (See Figure 5.)

Figure 5 Layer 2 Dispatch Mode

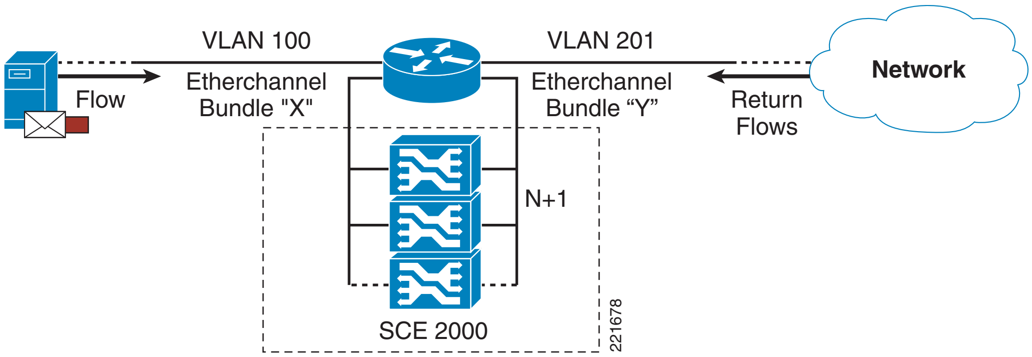

EtherChannel bundles need to be created to process both the host-side VLAN traffic and the network-side VLAN traffic. The SCE performs its packet inspection and configured policy enforcement, and returns the packet to the original data path. Any single flow must be dispatched to the same SCE to maintain state information. By default, the Layer 2 switching and EtherChannel hash algorithm balances the flows based on the IP header information, and the resulting path selection is the same for all packets in a particular flow. Because most customers deploy redundant cores and aggregation layers, it is worth mentioning how asymmetric flows are handled by the SCE. Because the EtherChannel hash is predictable given a consistent set of inputs, cabling the SCE devices in a uniform fashion to redundant 6500/7600 chassis results in the flows being directed to the same SCE unit, regardless of which chassis the packet traverses.

VLAN translation is then used to avoid misdirecting Layer 2 switched packets back on an EtherChannel. The EtherChannel interface is configured as a trunk port, the trunk port carries the VLAN ID in the .1Q header payload and the .1Q VLAN ID is automatically rewritten at the network EtherChannel connection or by the SCE device from the User VLAN ID to the network EtherChannel VLAN ID. Without the use of VLAN translation, the packet is automatically dropped by the 6500/7600 line card port ASICs.

Configuration Example

In the configuration of the MGSCP solution, remember to enable load balancing on both the ingress and egress ports for the traffic going through the EtherChannel. Before 12.2(18) SXF, you can use the per-source/destination IP load balancing from the global command line. The load balancing hash should be performed on the Supervisor CPU using the following command:

```
<port-channel load-balance src-ip>
<port-channel load-balance dst-ip>
```

With releases after 12.2(18) SXF, if you are using line cards with DCEF capabilities, you can now do per-line card load balancing using the following command:

```
<port-channel load-balance dst-ip module <1-9>>
```

The following examples show the actual switch configuration for MGSCP for a Layer 2 MGSCP configuration.

Global Configuration

```
!
interface Port-channel10
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 201
 switchport mode trunk
 no ip address
 lacp max-bundle 2
!
interface Port-channel20
 switchport
 switchport trunk encapsulation dot1q
switchport trunk allowed vlan 100
 switchport mode trunk
 no ip address
 lacp max-bundle 2
!
!
port-channel per-module load-balance
port-channel load-balance dst-ip
port-channel load-balance src-ip module 1
!
vlan 100
 name Host-side
!
vlan 201
 name Network-side
```

Host Side Configuration

- Port configuration—Host side:

  ```
  !
  interface GigabitEthernet1/1
   switchport
   switchport access vlan 100
   switchport mode access
   no ip address
   load-interval 30
  !
  ```

- Port configuration—Host side:

  ```
  !
  interface GigabitEthernet2/13
   description <<< Connected to Backup host side >>>
   switchport
   switchport trunk encapsulation dot1q
   switchport trunk allowed vlan 100
   switchport mode trunk
   switchport vlan mapping enable
   switchport vlan mapping 201 100
   no ip address
   no cdp enable
   channel-protocol lacp
   channel-group 20 mode active
   lacp port-priority 65535
  !
  interface GigabitEthernet2/14
   description <<< Connected to SCE-2 host side >>>
   switchport
   switchport trunk encapsulation dot1q
   switchport trunk allowed vlan 100
   switchport mode trunk
   switchport vlan mapping enable
   switchport vlan mapping 201 100
   no ip address
   channel-protocol lacp
   channel-group 20 mode active
  !
  interface GigabitEthernet2/15
  ```

```
description <<< Connected to SCE-1 host side >>>
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 100
 switchport mode trunk
 switchport vlan mapping enable
 switchport vlan mapping 201 100
 no ip address
 channel-protocol lacp
 channel-group 20 mode active
!
```

Network Side Configuration

- Port configuration—Net side:

```
!
interface GigabitEthernet5/1
 switchport
 switchport access vlan 201
 switchport mode access
 no ip address
 load-interval 30
!
```

- Port configuration—Net side:

```
!
interface GigabitEthernet1/13
 description <<< Connected to Backup Net >>>
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 201
 switchport mode trunk
 switchport vlan mapping enable
 switchport vlan mapping 100 201
 no ip address
 no cdp enable
 channel-protocol lacp
 channel-group 10 mode active
 lacp port-priority 65535
!
interface GigabitEthernet1/14
 description <<< Connected to SCE-2 Net >>>
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 201
 switchport mode trunk
 switchport vlan mapping enable
 switchport vlan mapping 100 201
 no ip address
 channel-protocol lacp
 channel-group 10 mode active
!
interface GigabitEthernet1/15
description <<< Connected to SCE-1 Net >>>
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 201
 switchport mode trunk
 switchport vlan mapping enable
 switchport vlan mapping 100 201
 no ip address
 channel-protocol lacp
 channel-group 10 mode active
```

!

MGSCP Layer 2/Layer 3 Dispatch Mode

When deploying the MGSCP solution in a mixed Layer 2/Layer 3 environment; for example, at the aggregation layer where the host-facing ports are Layer 2 and the core-facing ports are Layer 3, the only change needed is for a switched virtual interface (SVI) to allow packets to be routed in and out of the network EtherChannel. (See Figure 6.)

Figure 6 Layer 2/Layer 3 Deployment

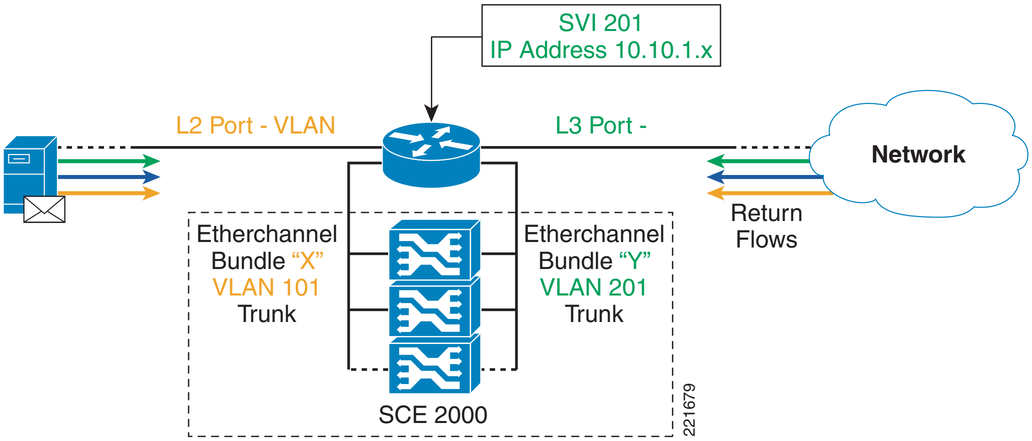

MGSCP Layer 3 Dispatch Mode

In the event that MGSCP needs to be deployed in an all-Layer 3 environment; for example, at the core layer of the data center or at the aggregation layer when running a completely routed data center design, you need to use an ACL on the 6500/7600 to define an action for a range of IP addresses. The EtherChannel is now configured as Layer 3 with an assigned IP address, but the hashing algorithm still occurs based on the flow IP address. (See Figure 7.)

Figure 7 Complete Layer 3 Deployment

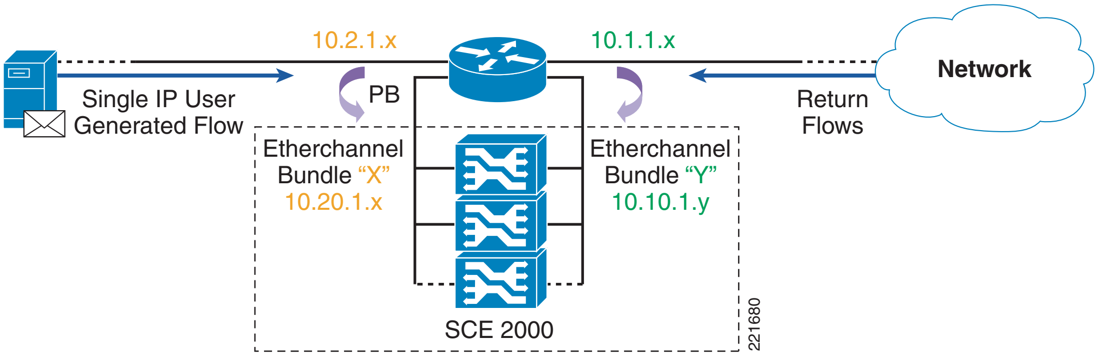

N+1 Redundancy

The MGSCP solution takes advantage of the network infrastructure to provide for an N+1 level of failover support. EtherChannel links can be monitored using Link Aggregation Control Protocol (LACP). LACP allows for some links in the bundle to be active while others are waiting in standby mode. When an active link fails, one of the standby links is moved into the active list. Because the SCE devices are "bump in the wire", if a link on one side of the SCE fails, the opposite link must be brought down as well. The SCE devices can be configured with a "link reflection" for sub-second recognition. LACP also has a fast timeout configuration to detect a failed link in a bundle. This detection can occur within 1–2 seconds. (See Figure 8.)

Figure 8 Link Failure

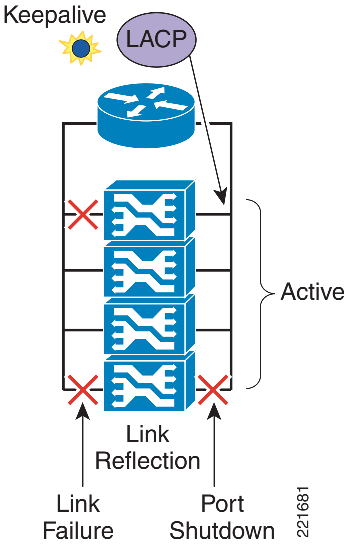

SCE Management and Policy Creation

The SCE solution relies heavily on the management stations for configuration and day-to-day operation. The Engage Console becomes the primary operational window that operators use to control network policy and to generate reports, as well as to configure SCE units operations. (See Figure 9.)

Figure 9 SCAS BB Console Window

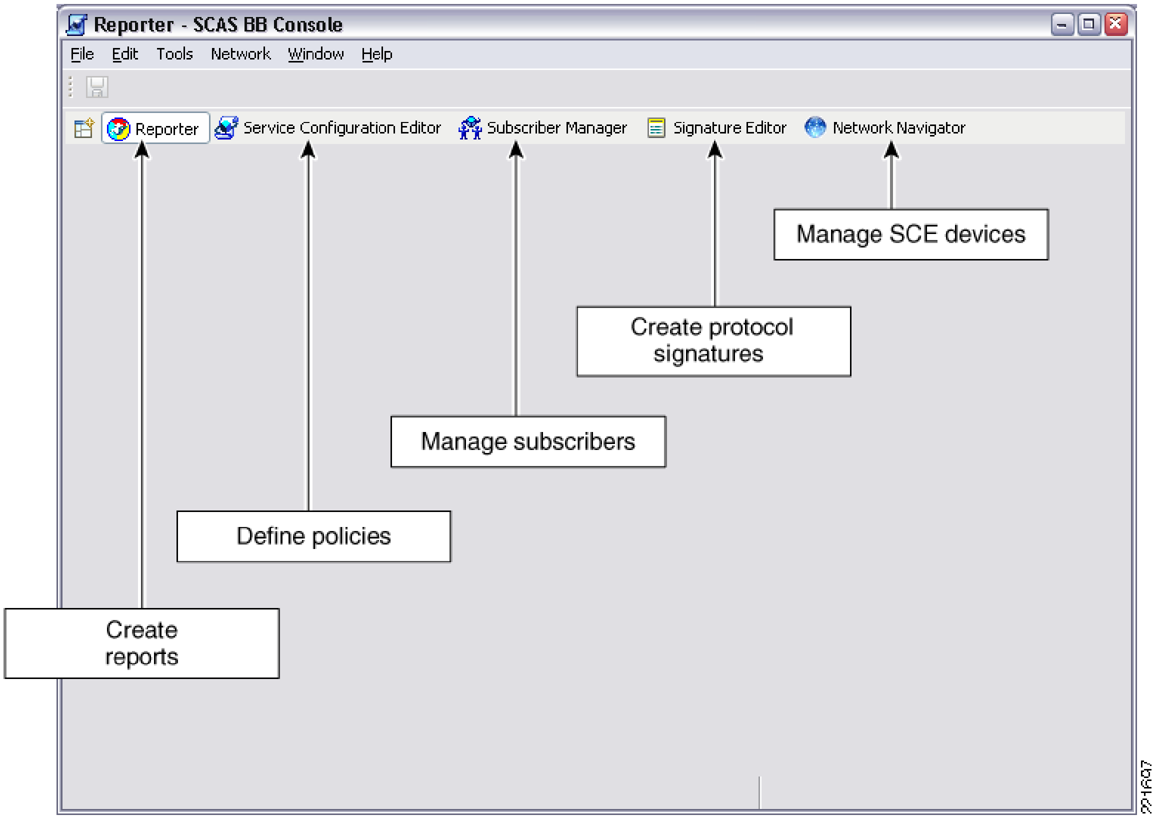

For example, the Network Navigator tab allows the operator to see the SCE units in operation and to group them in a logical collection. (See Figure 10.)

This is also where policies can be deployed to the SCE units; either to a single SCE device, or to a collection of SCE units. It may be useful for an operator to configure a policy and deploy it to one SCE device. Then, upon checking for issues such as performance impact and desired behavior, they can deploy the policy to the rest of the SCEs in the network.

Figure 10 Network Navigator Tab

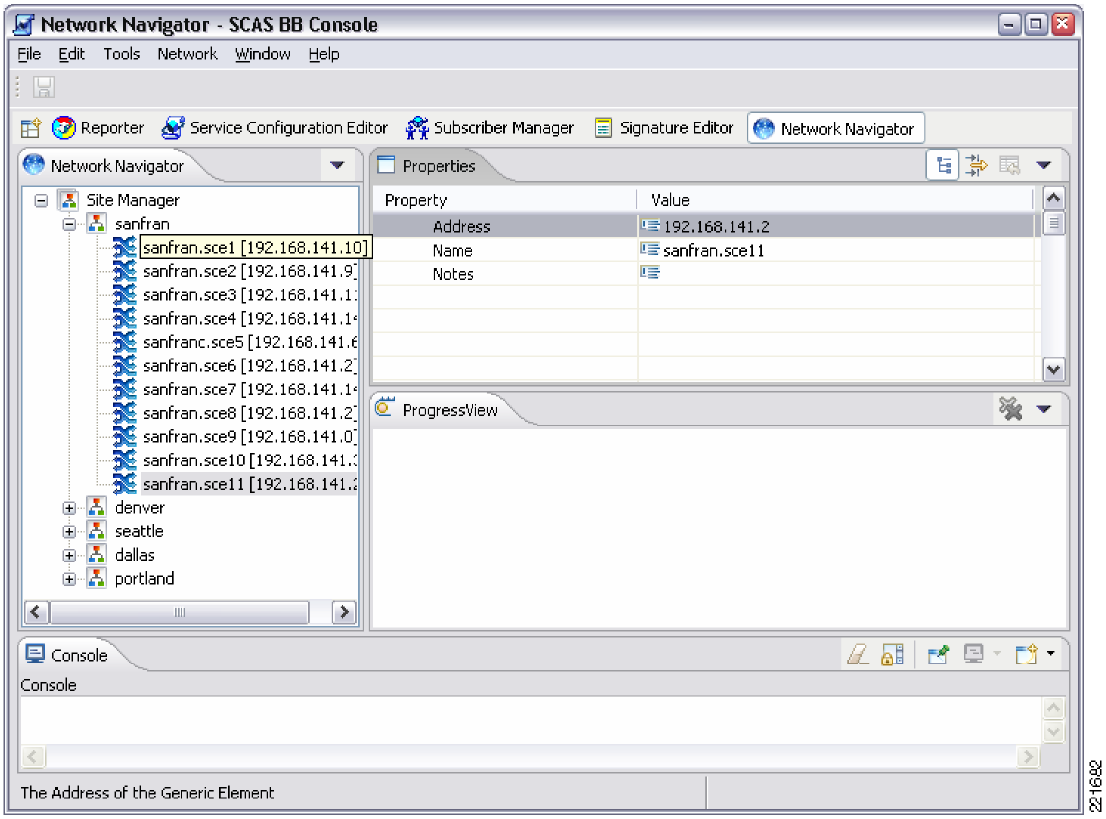

Network Navigator - SCAS BB Console
File Edit Tools Network Window Help
Reporter
Service Configuration Editor
Subscriber Manager
Signature Editor
Network Navigator
Network Navigator
Site Manager
sanfran
sanfran.sce1 [192.168.141.10]
sanfran.sce2 [192.168.141.9
sanfran.sce3 [192.168.141.1
sanfran.sce4 [192.168.141.1
sanfranc.sce5 [192.168.141.6
sanfran.sce6 [192.168.141.2
sanfran.sce7 [192.168.141.1
sanfran.sce8 [192.168.141.2
sanfran.sce9 [192.168.141.0
sanfran.sce10 [192.168.141.
sanfran.sce11 [192.168.141.2
denver
seattle
dallas
portland
Properties
ProgressView
Console
Console
221682

Property	Value
Address	192.168.141.2
Name	sanfran.sce11
Notes	

The Address of the Generic Element

Deploying a New SCE

This section outlines the steps necessary to install a new SCE into the network and bring it into production.

Note A console connection is required.

Step 1 Run the setup script, and enter basic configuration information such as IP address, netmask, SNMP strings, NTP, and DNS. (See Figure 11.)

Figure 11 *Basic Configuration of SCE*

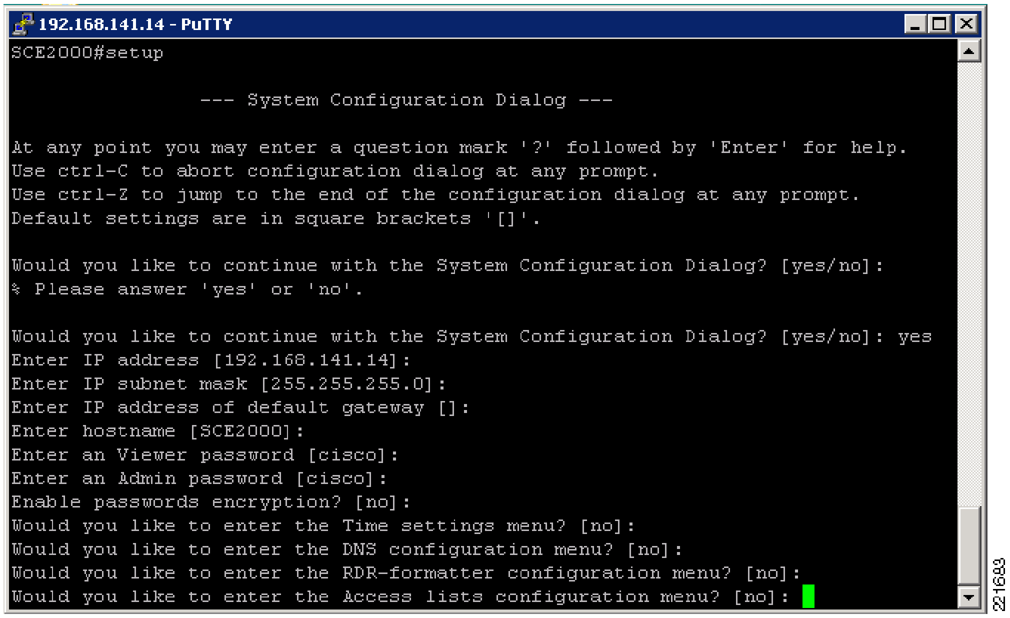

```
SCE2000#setup

                --- System Configuration Dialog ---

At any point you may enter a question mark '?' followed by 'Enter' for help.
Use ctrl-C to abort configuration dialog at any prompt.
Use ctrl-Z to jump to the end of the configuration dialog at any prompt.
Default settings are in square brackets '[]'.

Would you like to continue with the System Configuration Dialog? [yes/no]:
% Please answer 'yes' or 'no'.

Would you like to continue with the System Configuration Dialog? [yes/no]: yes
Enter IP address [192.168.141.14]:
Enter IP subnet mask [255.255.255.0]:
Enter IP address of default gateway []:
Enter hostname [SCE2000]:
Enter an Viewer password [cisco]:
Enter an Admin password [cisco]:
Enable passwords encryption? [no]:
Would you like to enter the Time settings menu? [no]:
Would you like to enter the DNS configuration menu? [no]:
Would you like to enter the RDR-formatter configuration menu? [no]:
Would you like to enter the Access lists configuration menu? [no]:
```

Step 2 Connect the cabling of the device to the network. (See Figure 12.)

Ensure that you have provisioned the network completely, as previously outlined.

Figure 12 *Cable the SCE to the Network*

Step 3 Using the SCAS BB console, add the SCE into the Network Navigator tab. (See Figure 13.)

You need to provide the IP address and the name of the device.

Figure 13 Adding the SCE to the Network Navigator

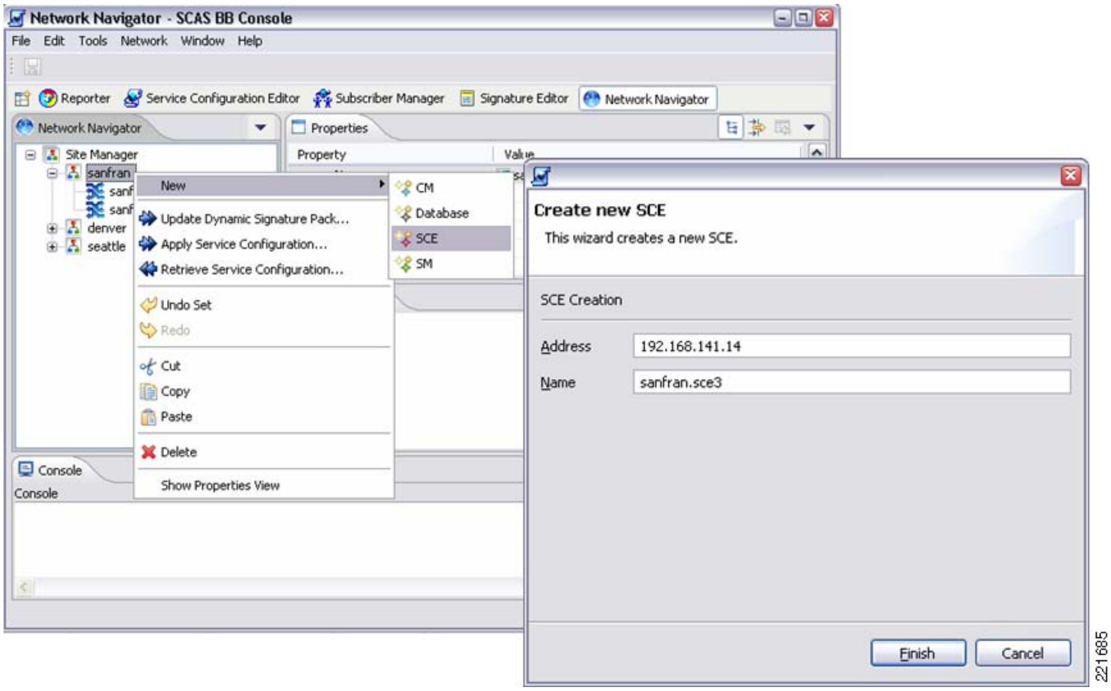

Create new SCE

This wizard creates a new SCE.

SCE Creation

Address 192.168.141.14

Name sanfran.sce3

Step 4 To ensure that the device is ready for policy deployment, it is important to ensure that the latest firmware has been installed on the device. (See Figure 14.)

Figure 14 Downloading Latest Firmware to the New SCE

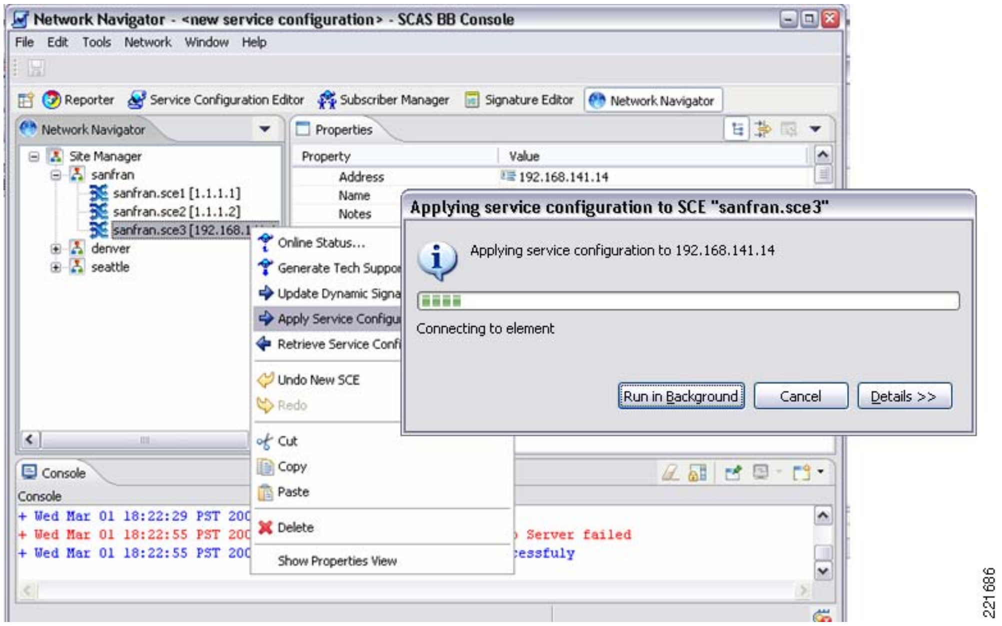

Property	Value
Address	192.168.141.14
Name	
Notes	

Applying service configuration to SCE "sanfran.sce3"

Applying service configuration to 192.168.141.14

Connecting to element

```
+ Wed Mar 01 18:22:29 PST 200
+ Wed Mar 01 18:22:55 PST 200                              Server failed
+ Wed Mar 01 18:22:55 PST 200                              essfuly
```

Step 5 The final step is to deploy any custom policies that have been created to the unit. (See Figure 15.)

When this step is complete, the unit is now in sync with any other SCE engines deployed in the cluster.

Figure 15 Deploy Site Policy Configuration (PQB)

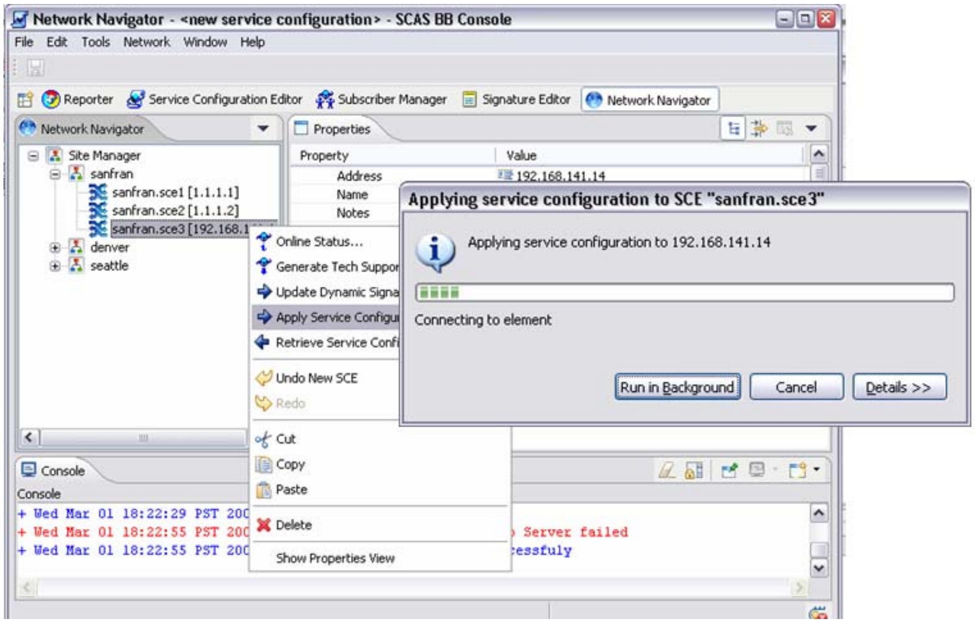

Applying service configuration to SCE "sanfran.sce3"

Applying service configuration to 192.168.141.14

Connecting to element

```
+ Wed Mar 01 18:22:29 PST 200
+ Wed Mar 01 18:22:55 PST 200                                   Server failed
+ Wed Mar 01 18:22:55 PST 200                                   essfuly
```

Changing a Policy

As a routine course of operation, the policies in place may need to be adjusted for things such as new applications, higher or lower bandwidth demands, and so on. The steps to edit a policy are as follows.

Step 1 In this example, the global service policy is changed to limit P2P traffic from 30 percent to 25 percent. This is accomplished by first opening the Policy Editor in the SCAS BB console. (See Figure 16.)

You can adjust the bandwidth in a uni-directional fashion so care should be taken to ensure you understand which traffic flow you want to impact.

Figure 16 Policy Editor Window

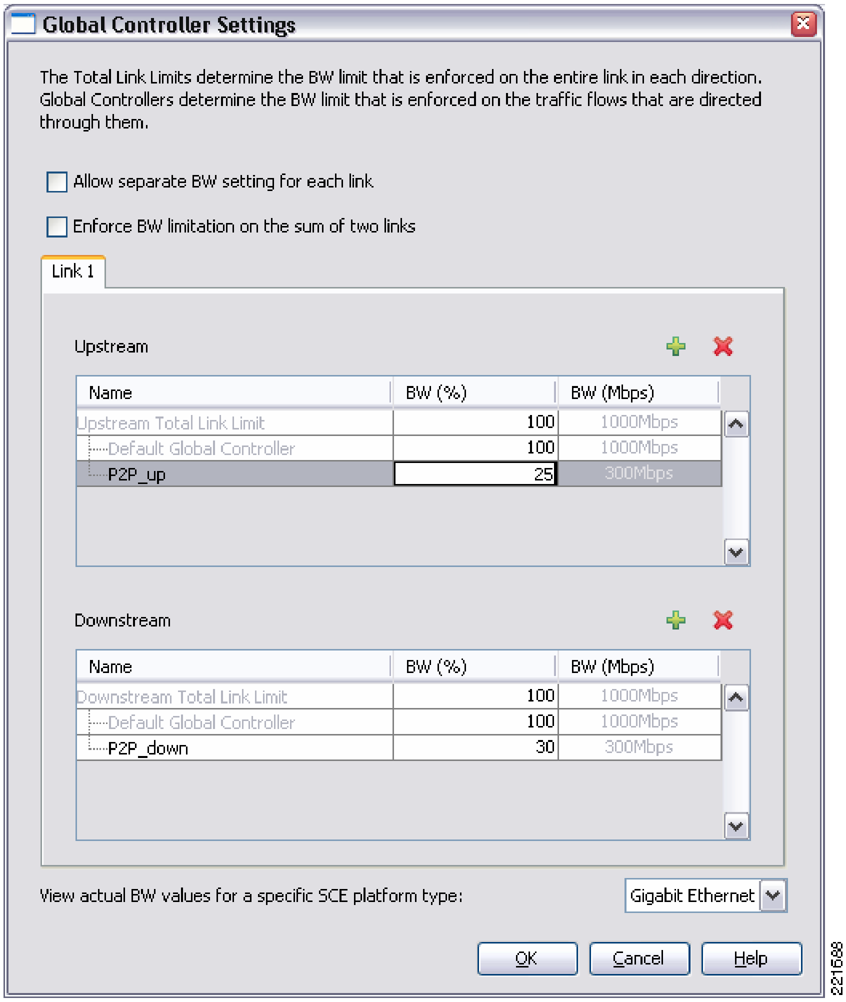

Global Controller Settings

The Total Link Limits determine the BW limit that is enforced on the entire link in each direction. Global Controllers determine the BW limit that is enforced on the traffic flows that are directed through them.

□ Allow separate BW setting for each link

□ Enforce BW limitation on the sum of two links

Link 1

Upstream

Name	BW (%)	BW (Mbps)
Upstream Total Link Limit	100	1000Mbps
Default Global Controller	100	1000Mbps
P2P_up	25	300Mbps

Downstream

Name	BW (%)	BW (Mbps)
Downstream Total Link Limit	100	1000Mbps
Default Global Controller	100	1000Mbps
P2P_down	30	300Mbps

View actual BW values for a specific SCE platform type: Gigabit Ethernet

Step 2 After the policy is updated, select the SCE devices to which this policy is to be downloaded; either an individual SCE or a group. (See Figure 17.)

Figure 17 Applying Policy to the SCE

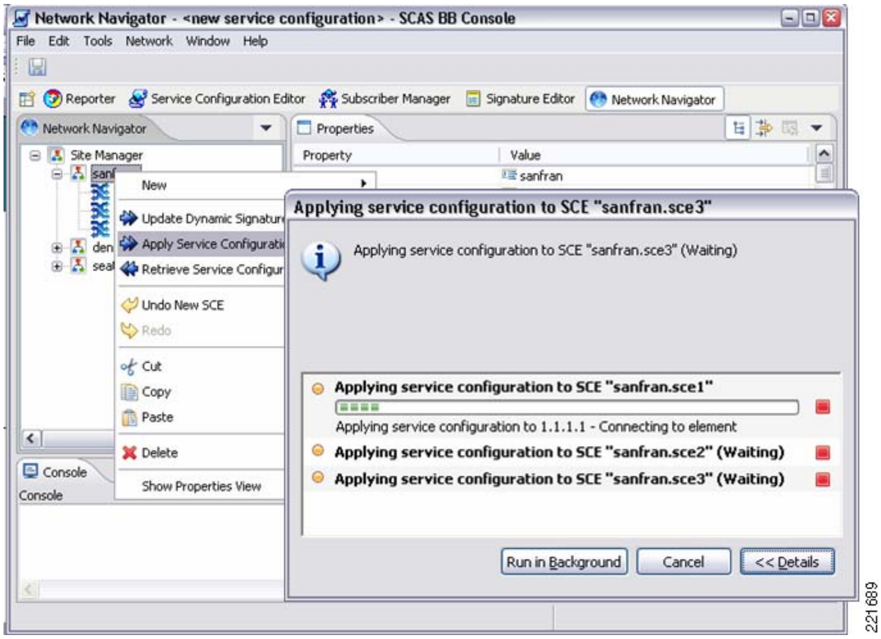

Applying service configuration to SCE "sanfran.sce3"

Applying service configuration to SCE "sanfran.sce3" (Waiting)

Applying service configuration to SCE "sanfran.sce1"

Applying service configuration to 1.1.1.1 - Connecting to element

Applying service configuration to SCE "sanfran.sce2" (Waiting)

Applying service configuration to SCE "sanfran.sce3" (Waiting)

Updating Signatures

This section deals with the routine updating of the signature files. Cisco publishes protocol packs that customers can download. Cisco recommends that customers use the latest release of the protocol pack because these updates provide greater application awareness.

After downloading the protocol pack, the network operator uses the SCAS BB Console to select the SCE(s) to which to deploy the protocol pack. The updates can then be installed. (See Figure 18.)

It is useful to keep all the SCE(s) in one location or cluster on the same protocol pack.

Figure 18 Updating the Signature Pack

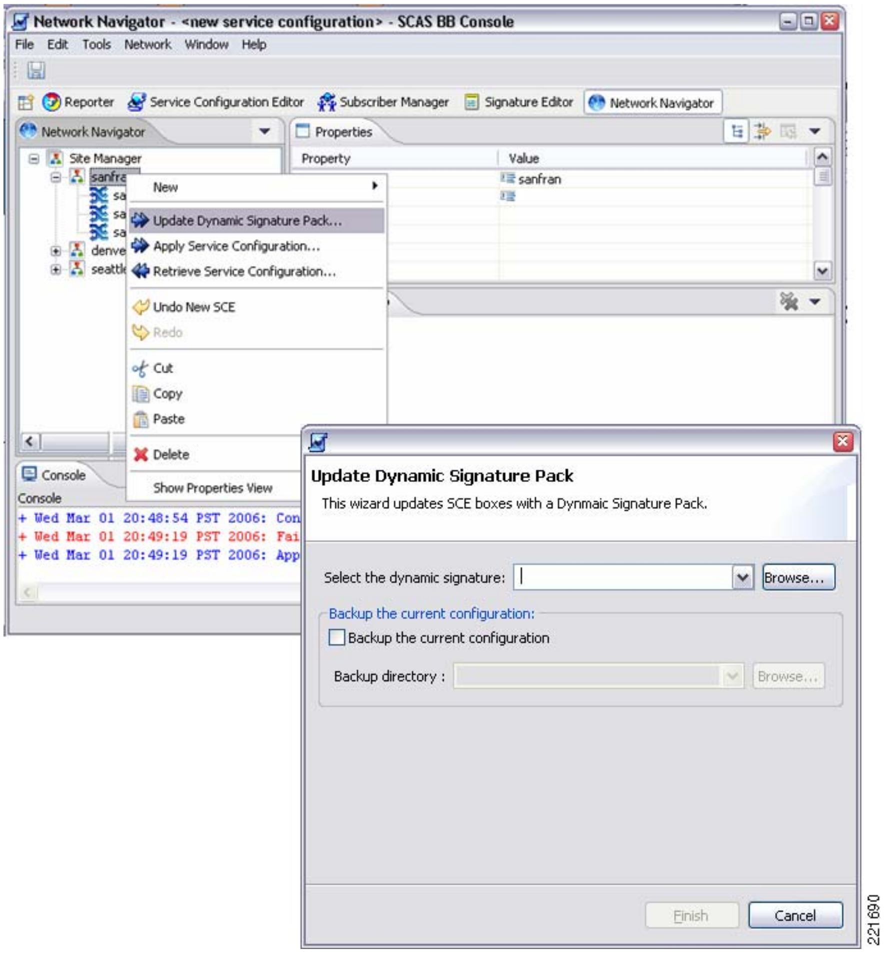

```
+ Wed Mar 01 20:48:54 PST 2006: Con
+ Wed Mar 01 20:49:19 PST 2006: Fai
+ Wed Mar 01 20:49:19 PST 2006: App
```

Update Dynamic Signature Pack

This wizard updates SCE boxes with a Dynmaic Signature Pack.

Select the dynamic signature:

Backup the current configuration:

□ Backup the current configuration

Backup directory :

Reports

The ability to collect, analyze, and report are a key feature of the SCE product. Following is a list of many of the reports that are part of the Service Control solution. Users can customize existing reports or create new report types. Through the use of the protocol packs that Cisco distributes routinely, the SCE solution stays current with new signature files and port-based applications. Users can again define their own signatures, but the product has built-in support for 100+ protocols.

The following are examples of the reports generated by the Service Control solution.

- General reports
  - Both direction bandwidth per service
  - Relative consumption of top 20 subscribers
  - Browsing counters service popularity
  - Services for a given subscriber
  - Top protocols by total volume
  - Global aggregated usage per service

  - Global concurrent session per service
  - Hourly top 20 subscribers
  - Global hourly usage volume downstream per service
- Voice reports
  - Bi-directional bandwidth per VoIP service
  - Global concurrent session per VoIP service
  - Global hourly call minutes per VoIP service
  - Hourly SIP top talkers
  - Top SIP domains by volume
- Streaming reports
  - Top streaming hosts by volume
  - Top RSTP hosts ordered by volume
- P2P reports
  - Subscribers using P2P
  - Top P2P protocols by total volume
  - Top P2P downloaders
- E-mail and news reports
  - Top newsgroups by volume
  - Top newsgroups by hit count
  - Top e-mail recipients by volume
  - Top e-mail senders by volume
- Security reports
  - Global DoS rate
  - Global scan attack rate
  - Infected subscribers
  - Top scanning/attacking host
  - Top DoS attacked host
  - Top DoS attacked subscribers

Example reports are shown in the following figures.

Figure 19 Global Bandwidth per Service

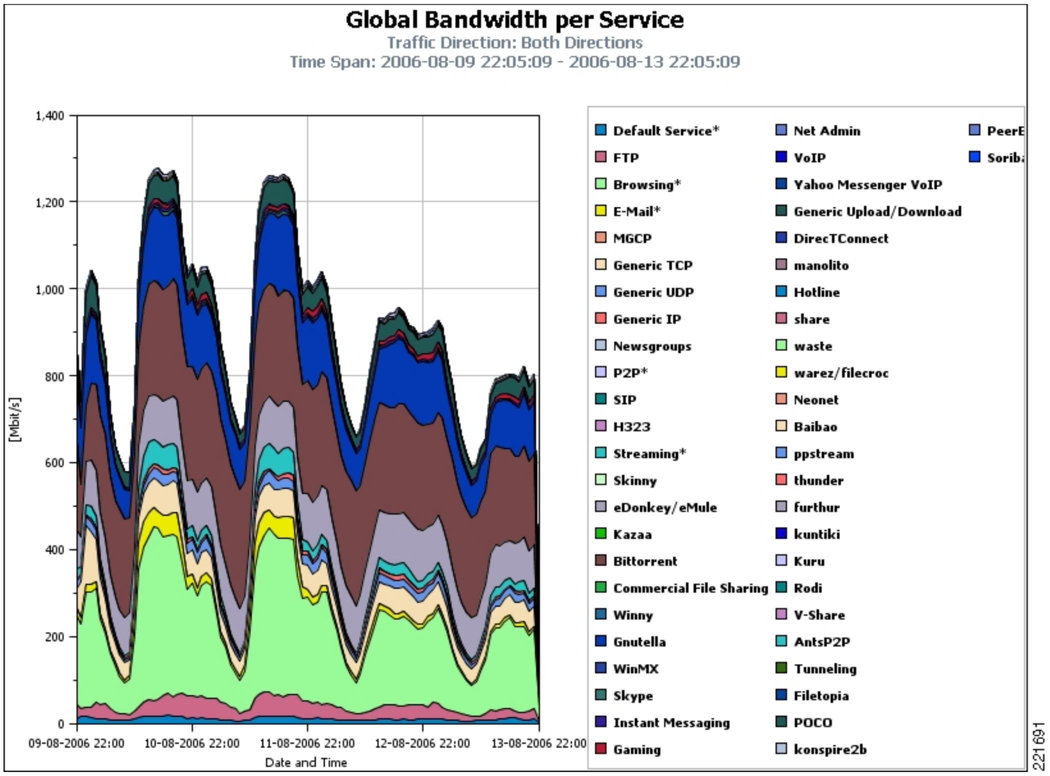

Figure 20 Global Bandwidth per Service—Downstream Direction

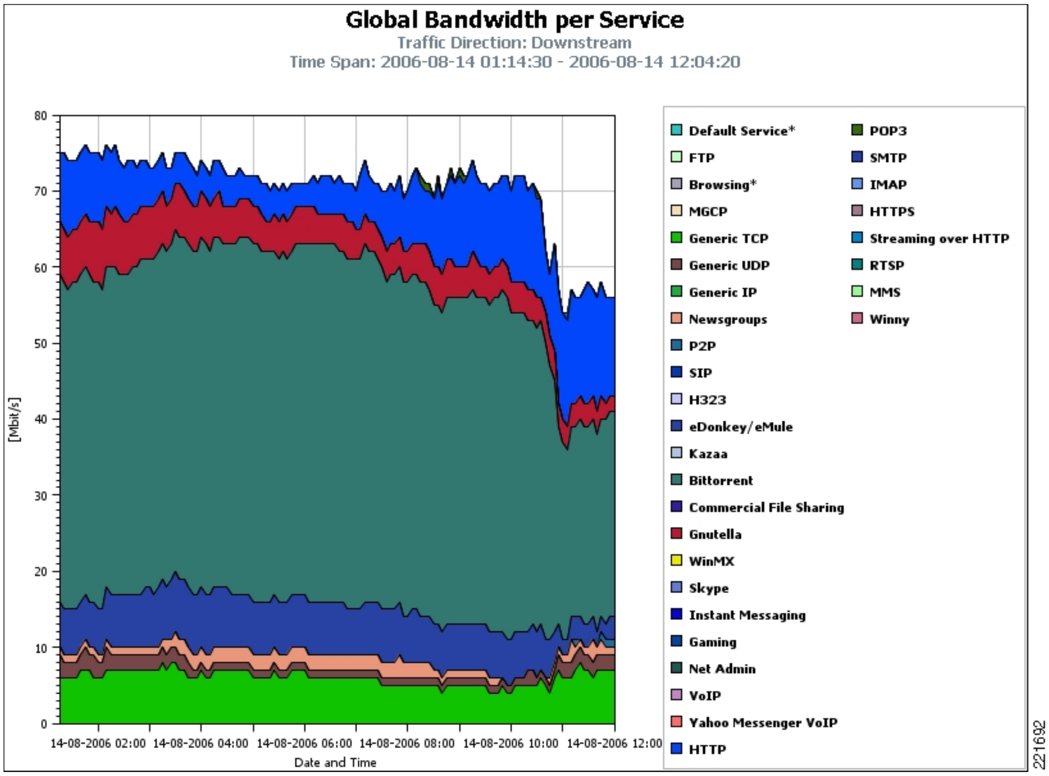

Figure 21 Daily Usage Volume

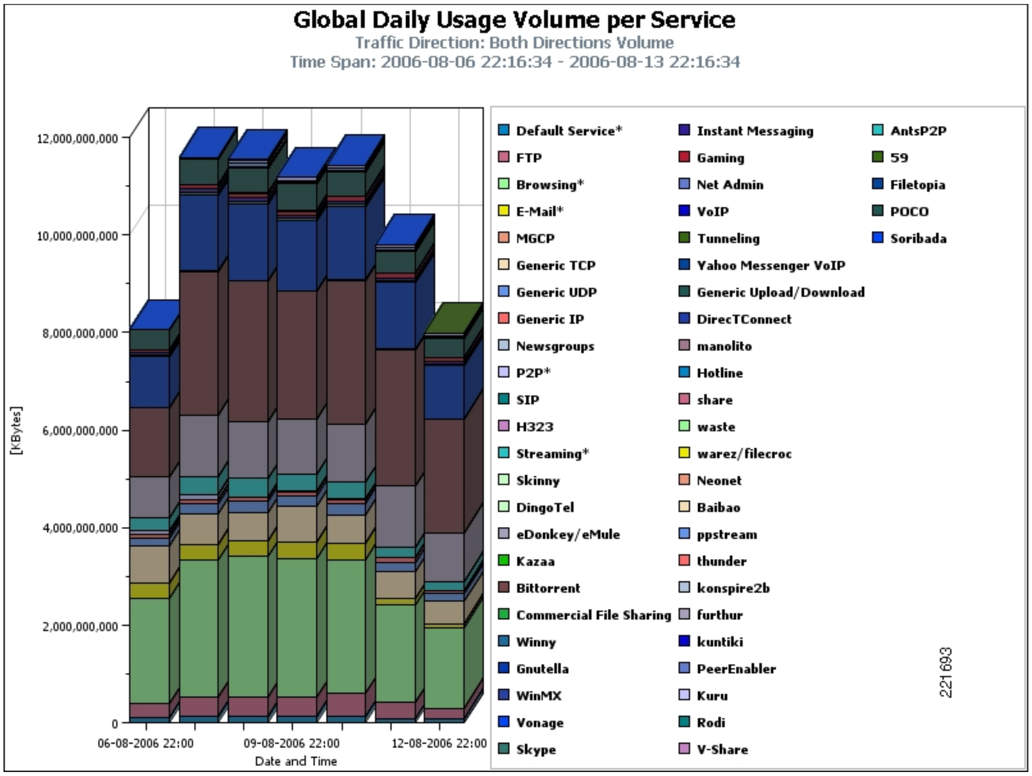

Figure 22 Top Subscribers in Selected Service

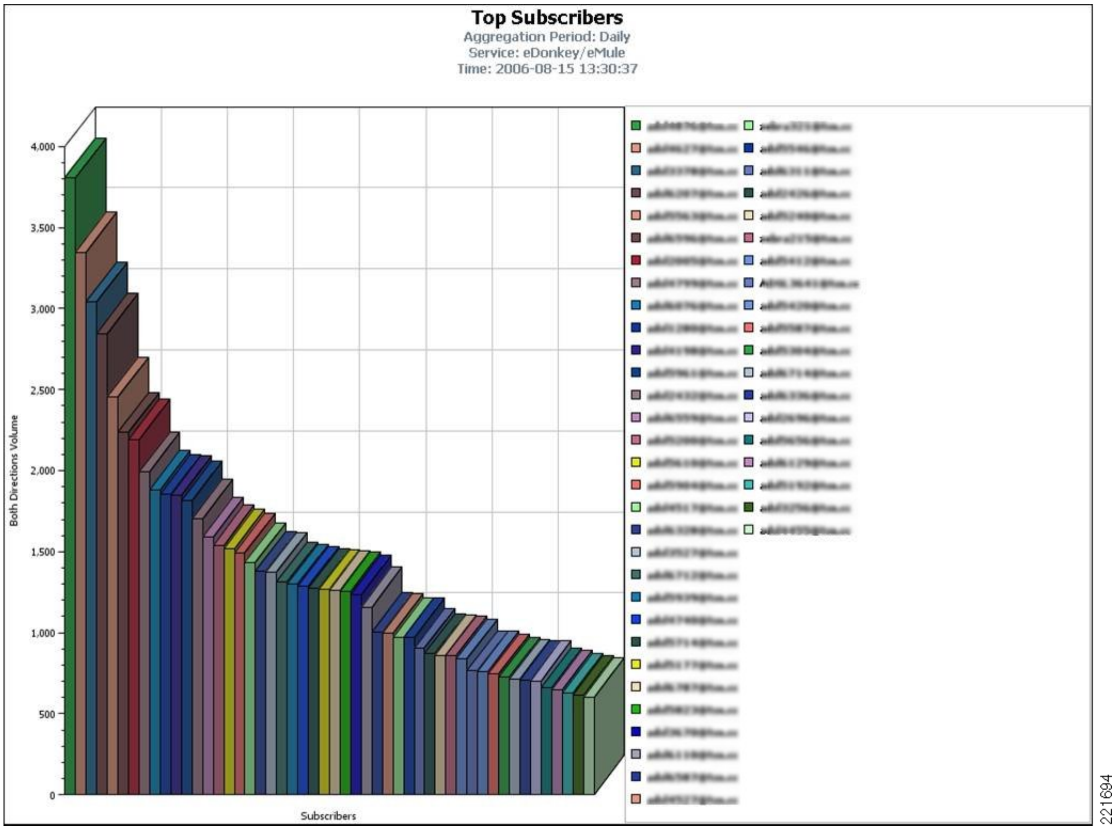

Figure 23 Subscriber Usage Bandwidth

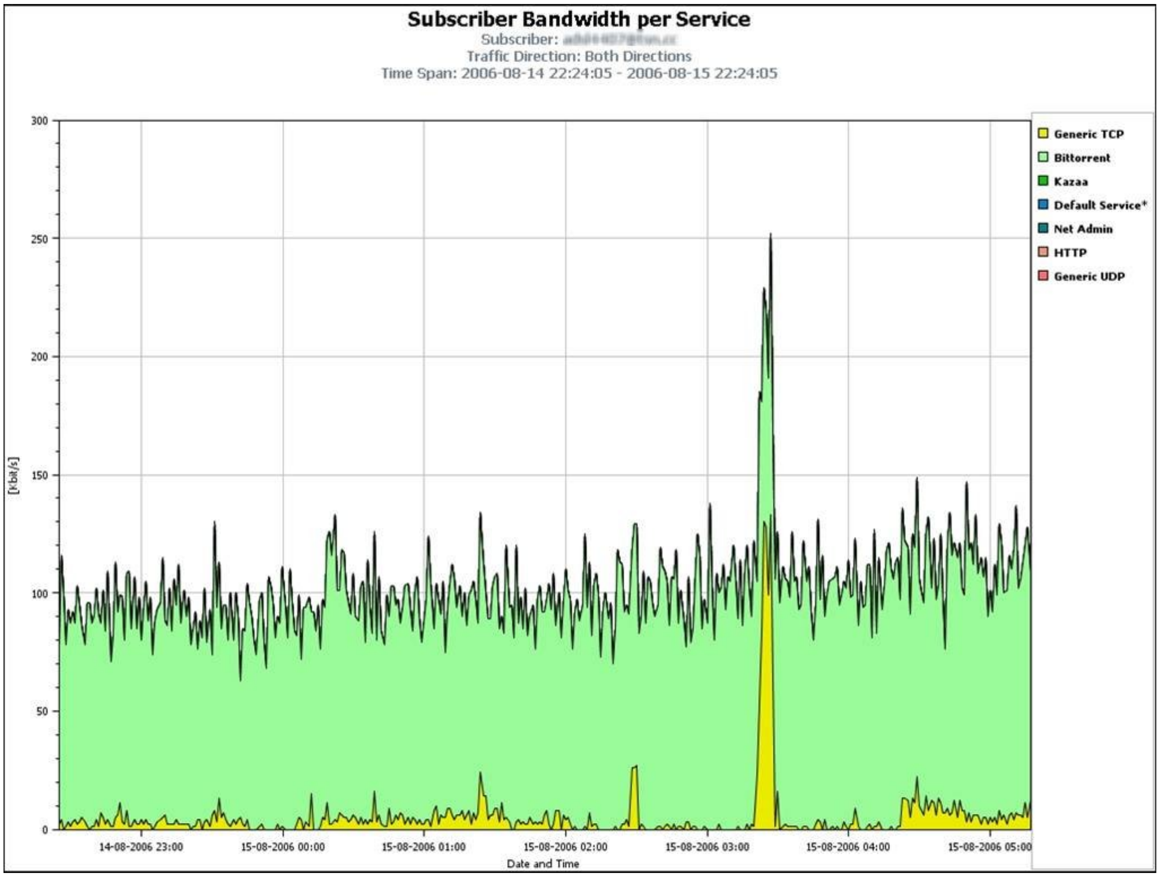

Figure 24 VoIP MOS

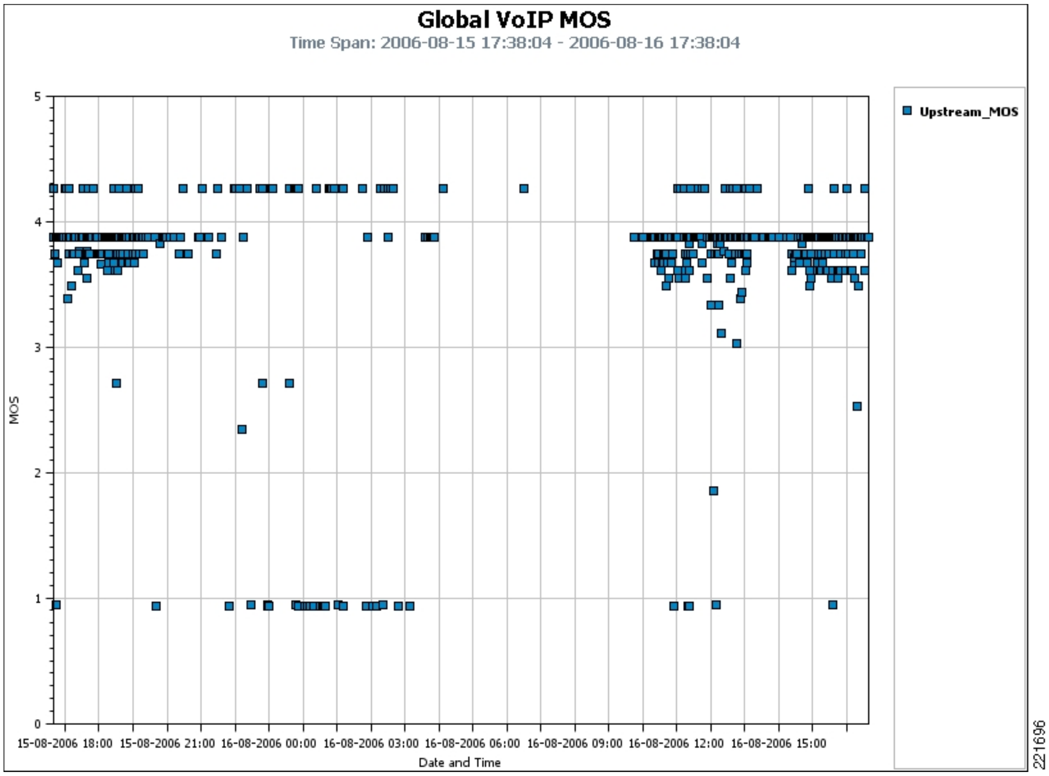

Appendix A—SCE and PISA

Recently, Supervisor 32 with PISA hardware was introduced. This section provides a comparison of the SCE and PISA products.

At this time, testing of the Supervisor 32 with PISA has not been performed, but may be undertaken and documented in the future.

Supervisor Engine 32 PISA is primarily targeted at the enterprise in the WAN edge:

- Embedded IOS-based DPI and IP services
- Application awareness and intelligent classification with NBAR
- Rapid security protection with flexible packet matching
- Full integration with other Supervisor 32 features such as IPv4/IPv6 in hardware, advanced multicast and MPLS, enhanced manageability, HA with NSF/SSO, and more.

SCE had traditionally been positioned for the service provider market by offering the following:

- Hardware-accelerated DPI
- Application and subscriber-awareness to manage traffic by end-user and Layer 7 applications
- Visibility, P2P management, malicious traffic detection, and service creation

The ability to cluster the SCE and provide a multi-gigabit capability, however, make this product suited for inclusion in the data center as a method for inspecting traffic between the campus network and the data center, or in the high speed Metro Ethernet connections that can exist between multi-site data center environments.

The following tables provide a comparison of the two technologies.

	Sup32-PISA	**SCE**
Primary target market	Enterprise (WAN edge)	Service provider edge and data center
Deep packet inspection	Yes	Yes
Routing/switching feature set	Yes (IPv4/6 in HW, advanced multicast and MPLS, and so on)	No
Subscriber-awareness	No	Yes
Deployment model	Embedded on Catalyst 6500 Supervisor	Appliance (SCE1010, SCE2020)
OS integration	IOS integrated	SEF component (ISG integration, BPM integration, uBR-CMTS awareness…)
Management	Embedded + GUI (CLI, EEM, CSM, QPM, SNMP, and so on)	Service Control toolset (CM, SM, Console GUI, and so on)
Maximum packet inspection depth	8192 bytes	1600 bytes
Jumbo frame support	Yes	No
Performance	2Gbps DPI, 32Gbps System	2Gbps/4Gbps DPI
Maximum number of flows (bi-directional)	256K flows	1M flows
Support IPv6	Yes (no DPI)	No
MPLS/VPN	Yes (no DPI)	Yes (with DPI)
Multicast	Yes (no DPI)	No
Supports regular expressions	Yes	Yes
Actions supported	Police, permit, drop, mark, log	Police, permit, drop, mark, redirect
Port density	Based on integrated solution SUP32 (1,000+)	Based on bump-in-the-wire concept (4 port per SCE2020)

Number of pre-defined L7 protocols	90+	100+
Stateful	Yes	Yes
Interface support	Gige, POS, ATM, T1/E1, Channelized interfaces	Gige
Protocol customization	Yes (CLI)	Yes (GUI)
Transparent protocol update	Yes	Yes (monthly protocol packs on CCO)

	SUP32 PISA	SCE (1000/2000)
Classification		
P2P	BitTorrent, DirectConnect, KaZaa GNUtella, WinMX, eDonkey	Yes (extensive 30+)
Business protocols	CITRIX, SAP, EXCHANGE, SQL	No Layer 7 (As of Jan.07)
HTTP.URL Regexp	Yes	Yes (100,000 URL entries)
HTTP.Mime	Yes	Yes
Routing protocols	Yes	No (Customizable)
Management protocols (Telnet, SNMP, and so on)	Yes	Yes
Voice protocol	RTCP, SSCP, SIP, H323, Skinny, MGCP, RTP	SIP, H323, Skinny, MGCP, Skype, RTP, RTCP
Worm signatures	Yes (FPM, CLI and XML)	Yes (GUI Signature Editor)
Anomaly detection	No	Yes
DSCP	Yes	No
User/subscriber	No	Yes
ACL	Yes	No
Integrated content filtering engine	No	Yes
Bandwidth Management		
	Yes (per application/interface)	Yes (per application/user)

	SUP32 PISA	SCE (1000/2000)
Reporting		
Bandwidth per application	Yes	Yes

Subscriber (demographic reports)	No	Yes
Security reports	Yes	Yes
Syslog export	Yes	No
SNMP export	Yes	Yes
NetFlow export	Yes (not for DPI)	Future
Proprietary export protocol	No	RDR
SPAN, ERSPAN	Yes	No
Reporting tool	Yes (CSM, QPM, SNMP-tools)	Yes (Integrated CM and Console, SNMP tools)
Integrated Reporting Tool	No	Yes
Manageability		
CLI	Yes	Yes
GUI	Yes (CSM, QPM, SNMP tools)	Yes (Integrated Console, SNMP tools)
Embedded Event Manager	Yes	No
Integrated GUI	No	Yes
CiscoWorks integration	Yes	No
Customer API	Yes (XML)	Yes (Java)