

NOKIA

Traffic management in Voyager for IPSO 3.8NET

The product described in this document is still under development by Nokia Networks. However, in the interest of offering early possibility to our customers to evaluate the documentation, this documentation is provided in draft form. Therefore the customer understands that the information in this document is subject to change without notice and describes only the prototype product defined in the introduction of this documentation in its current state of development. Nokia Networks welcomes customer comments as part of the process of continuous development and improvement of its products and the documentation.

This document is not a final customer document and Nokia Networks does not take responsibility for any errors or omissions in this document. No part of it may be reproduced or transmitted in any form or means without the prior written permission of Nokia Networks. The document has been prepared to be used by professional and properly trained personnel, and the customer assumes full responsibility when using it.

The information or statements given in this document concerning the suitability, capacity, or performance of the mentioned hardware or software products cannot be considered binding but shall be defined in the agreement made between Nokia Networks and the customer.

Nokia Networks WILL NOT BE RESPONSIBLE IN ANY EVENT FOR ERRORS IN THIS DOCUMENT OR FOR ANY DAMAGES, INCIDENTAL OR CONSEQUENTIAL (INCLUDING MONETARY LOSSES), that might arise from the use of this document or the information in it. UNDER NO CIRCUMSTANCES SHALL NOKIA BE RESPONSIBLE FOR ANY LOSS OF USE, DATA, OR INCOME, COST OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, PROPERTY DAMAGE, PERSONAL INJURY OR ANY SPECIAL, INDIRECT, INCIDENTAL, PUNITIVE OR CONSEQUENTIAL DAMAGES HOWSOEVER CAUSED.

THE CONTENTS OF THIS DOCUMENT ARE PROVIDED "AS IS". EXCEPT AS REQUIRED BY APPLICABLE MANDATORY LAW, NO WARRANTIES OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT, ARE MADE IN RELATION TO THE ACCURACY, RELIABILITY OR CONTENTS OF THIS DOCUMENT. NOKIA RESERVES THE RIGHT TO REVISE THIS DOCUMENT OR WITHDRAW IT AT ANY TIME WITHOUT PRIOR NOTICE.

This document and the product it describes are protected by copyright according to the applicable laws.

NOKIA and Nokia Connecting People are registered trademarks of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective companies, and they are mentioned for identification purposes only.

Copyright © Nokia Corporation 2003. All rights reserved. Reproduction, transfer, distribution or storage of part or all of the contents in this document in any form without the prior written permission of Nokia is prohibited.

Contents

	Contents	3
1	Access control lists (ACL)	5
2	Creating access control lists	7
3	Deleting access control lists	9
4	Applying access control lists to interfaces	11
5	Removing access control lists from interfaces	13
6	Access control list (ACL) rules	15
7	Adding new rules to access control lists	17
8	Modifying access control list rules	19
9	Deleting access control list rules	23
10	Aggregation classes	25
11	Creating aggregation classes	27
12	Deleting aggregation classes	29
13	Associating aggregation classes with rules	31
14	Queue classes	33
15	Creating queue classes	35
16	Deleting queue classes	37
17	Setting or modifying queue class configuration values	39
18	Associating queue classes with interfaces	43
19	ATM QoS	45
20	Creating QoS descriptors	47
21	Deleting ATM QoS descriptors	49
22	Associating an ATM QoS descriptor with an interface and a virtual channel	51
23	Common open policy server (COPS)	53

24	Configuring COPS client IDs and policy decision points	55
25	Configuring security parameters for COPS client IDs	57
26	Assigning roles to specific interfaces	59
27	Activating and deactivating the COPS client	61
28	Changing client IDs associated with specific Diffserv configurations	63
29	Deleting client IDs	65
30	Example of rate shaping configuration	67
31	Example of expedited forwarding configuration	69

1

Access control lists (ACL)

The traffic management software allows packet streams to be filtered, shaped, or prioritized. The prioritisation mechanisms conform to RFC 2598 and RFC 2597 of the IETF.

Traffic is separated into discrete streams, or classified, through an access control list (ACL). Traffic is metered to conform to throughput goals with an aggregation class (AGC). The combination of these control blocks form the basis of the filtering, shaping, and prioritization tools. A queue class is used to implement an output scheduling discipline to prioritize traffic.

Logically, the ACLs and the AGCs are placed inline to the forwarding path. You can configure ACLs and AGCs to process all incoming traffic from one or more interfaces, or to process all outgoing traffic from one or more interfaces. IPSO supports ACLs for both IPv4 and IPv6 traffic.

Packet filtering description

Traffic that is classified can be filtered immediately. The actions for filtering are:

- Accept—The accept action forwards the traffic.
- Drop—The drop action drops the traffic without any notification.
- Reject—The reject action drops the traffic and sends an ICMP error message to the source.

For information on how to configure a packet filter, see *Access control list rules*.

Traffic shaping description

Traffic that is classified can be shaped to a mean rate. The shaper is implemented using a token bucket algorithm; this means that you can configure a burstsize from which bursts can "borrow." Measured over longer time intervals, the traffic will be coerced to the configured mean rate. Over shorter intervals, traffic is allowed to burst to higher rates. This coercion is accomplished by adding delay to packets that must wait for more tokens to arrive in the bucket. When more bursts arrive than can be accommodated by the shaping queue, then that traffic is dropped. Both outgoing and incoming traffic streams can be shaped.

To configure a shaper, see *Access control list rules*. Select shape as the action for one or more rules. See *Creating aggregation classes* for information about creating AGC meters. You should associate the AGC with the shaping rule(s) of the ACL.

Traffic queuing description

Traffic that is classified by an access control list (ACL) rule can be given preferential treatment according to RFC 2597 and RFC 2598. Higher-priority traffic must be policed to prevent starvation of lower-priority service traffic. Traffic that conforms to the configured policing rate is marked with the differentiated services codepoint (DSCP). When such traffic is processed by the output queue scheduler, it receives favorable priority treatment.

Some traffic is generated by networking protocols. This traffic should be given the highest queuing priority; otherwise, the link may become unstable. For this reason, the queue class (QC) configuration provides an internetwork control queue by default; some locally sourced traffic is prioritized to use that queue.

Prioritization is only relevant for outgoing traffic. Incoming traffic is never prioritized.

Use the DSfield in the access control list (ACL) to set the value for marking traffic that matches a given ACL rule. The QueueSpec is used to map a flow with the output queue.

To configure traffic queuing, see *Access control list rules* for information about creating ACL rules. Choose prioritize as the action for one or more rules. Enter the appropriate values in the DSfield and QueueSpec edit boxes. See *Creating aggregation classes* for information about creating aggregation class meters. You should associate the AGC with the prioritize rule(s) of the ACL.

For more information, see *Creating access control lists* and *Applying access control lists to interfaces*

2

Creating access control lists

Purpose

To set up an access control list (ACL), you must configure the interface(s) with which you want to associate the ACL and the Bypass option. IPSO supports both the IPv4 and IPv6 protocols. To configure an interface, see *Applying access control lists to interfaces*. The Bypass option denotes that the entire packet stream flowing out of the selected interfaces should not be classified, policed, or marked. Instead, the output queue scheduler should use the supplied IP TOS as an output queue lookup. Use the Bypass option to circumvent the classifier and policer for selected interfaces.



Steps

1. **Click Config on the home page.**
2. **IPSO supports both the IPv4 and IPv6 protocols. Select either the IPv4 or IPv6 protocols.**
 - a. For IPv4 ACLs, click the *Access List Configuration* link under the **Traffic Management** section.
 - b. For IPv6 ACLs, click the *IPv6* link. This takes you to the IPv6 page. Click the *Access List Configuration* link under the **Traffic Management** section.
3. **Enter a name for the ACL in the Create A New Access List edit box. Click Apply.**

Expected outcome

The access control list name, **Delete** check box, and **Bypass this Access List** field appear.

4. **To make your changes permanent, click Save.**

Further information

See *Access Control Lists* , *Deleting access control lists* and *Applying access control lists to interfaces*

3

Deleting access control lists



Steps

1. **Click Config on the home page.**
2. **IPSO supports both the IPv4 and IPv6 protocols. Select either the IPv4 or IPv6 protocols.**
 - a. For IPv4 ACLs, click the *Access List Configuration* link under the **Traffic Management** section.
 - b. For IPv6 ACLs, click the *IPv6* link. This takes you to the IPv6 page. Click the *Access List Configuration* link under the **Traffic Management** section.
3. **Click the Delete check box next to the access control list you want to delete. Click Apply.**

Expected outcome

The access control list name disappears from the Access List Configuration page.

4. **To make your changes permanent, click Save.**

Further information

See *Access control lists (ACL)* and *Removing access control lists from interfaces*

4

Applying access control lists to interfaces

Purpose

Note

You can apply the same interface with the same direction to an IPv4 access control list and to an IPv6 access control list. You cannot apply the same interface with the same direction to more than one IPv4 access control list to more than one IPv6 access control list.

The same interface can be configured to both user ACL and COPS ACL. In this case the user ACL is the matched first, without the default rule, and then the COPS ACL.



Steps

1. **Click Config on the home page.**
2. **IPSO supports both the IPv4 and IPv6 protocols. Select either the IPv4 or IPv6 protocols.**
 - a. For IPv4 ACLs, click the *Access List Configuration* link under the **Traffic Management** section.
 - b. For IPv6 ACLs, click the *IPv6* link. This takes you to the IPv6 page. Click the *Access List Configuration* link under the **Traffic Management** section.
3. **Click the link for the appropriate access control list in the ACL Name field.**

Expected outcome

You are taken to the page for that access control list.

4. **Select the appropriate interface from the ADD Interfaces drop-down window.**
 5. **Select either INPUT or OUTPUT from the Direction drop-down window. Click Apply.**
 6. **To make your changes permanent, click Save.**
-

Note

Selecting the "input" direction for an access control list with a rule whose action is set to "prioritize" is equivalent to setting the action to "skip."

Expected outcome

The new interface appears in the **Selected Interfaces** section.

Further information

Note

Only the default rule appears in the access control list until you create your own rule.

Further information

See *Access Control Lists* and *Removing access control lists from interfaces*

5

Removing access control lists from interfaces



Steps

1. Click **Config** on the home page.
2. **IPSO supports both the IPv4 and IPv6 protocols. Select either the IPv4 or IPv6 protocols.**
 - a. For IPv4 ACLs, click the *Access List Configuration* link under the **Traffic Management** section.
 - b. For IPv6 ACLs, click the *IPv6* link. This takes you to the IPv6 page. Click the *Access List Configuration* link under the **Traffic Management** section.
3. **Click the link for the appropriate access control list in the ACL Name field.**

Expected outcome

You are taken to the page for that access control list.

4. **Click the Delete check box next to the interface (to the right) under the Selected Interfaces section that you want to remove. Click Apply.**

Expected outcome

The interface disappears from the **Selected Interfaces** section.

5. **To make your changes permanent, click Save.**

Further information

See *Access control lists (ACL)*, *Applying access control lists to interfaces* and *Deleting access control lists*

6

Access control list (ACL) rules

An access control list (ACL) is a container for a set of rules, and traffic is separated into packet streams by the access control list. The content and ordering of the rules is critical. As packets are passed to an ACL, the packet headers are compared against data in the rule in a top-down fashion. When a match is found, the action associated with that rule is taken, with no further scanning done for that packet.

The following actions can be associated with a rule that is configured to perform packet filtering:

- Accept
- Drop
- Reject

The following additional actions can also be associated with a rule:

- Skip—skip this rule and proceed to the next rule
- Prioritize—give this traffic stream preferential scheduling on output
- Shape—coerce this traffic's throughput according to the set of parameters given by an aggregation class

Rules can be set up to match any of these properties:

- IP source address
- IP destination address
- IP protocol
- UDP/TCP source port
- UDP/TCP destination port

- TCP establishment flags—When selected, traffic matches this rule when it is part of the initial TCP handshake.
- Type of Service (TOS) for IPv4; Traffic Class for IPv6

The following values can be used to mark traffic:

- DiffServ codepoint (DSfield)
 - Queue Specifier (QueueSpec)
-

Note

The DSfield and QueueSpec field are used to mark and select the priority level.

Masks can be applied to most of these properties to allow wildcarding. The source and destination port properties can be edited only when the IP protocol is UDP, TCP, or the keyword "any."

All of these properties are used to match traffic. The packets that match a rule whose action is set to "prioritize" are marked with the corresponding DSfield and sent to the queue set by QueueSpec field. The DSfield and QueueSpec field can only be edited when the Action field is set to "prioritize."

For more information, see *Access control lists*, *Adding new rules to access control lists*, *Modifying access control list rules* and *Deleting access control list rules*

7

Adding new rules to access control lists



Steps

1. **Click Config on the home page.**
2. **IPSO supports both the IPv4 and IPv6 protocols. Select either the IPv4 or IPv6 protocols.**
 - a. For IPv4 ACLs, click the *Access List Configuration* link under the **Traffic Management** section.
 - b. For IPv6 ACLs, click the *IPv6* link. This takes you to the IPv6 page. Click the *Access List Configuration* link under the **Traffic Management** section.
3. **Click the link for the appropriate access control list in the ACL Name field.**

Expected outcome

You are taken to the page for that access control list.

4. **Click the Add New Rule Before check box. Click Apply.**

Expected outcome

This rule appears above the default rule.

Further information

After you create more rules, you can add rules before other rules. If you have four rules—rules 1,2,3, and 4—you can place a new rule between rules 2 and 3 by checking the **Add Rule Before** check box on rule 3.

5. **To make your changes permanent, click Save.**

Further information

See *Access control list (ACL) rules*, *Modifying access control list rules* and *Deleting access control list rules*

8

Modifying access control list rules



Steps

1. Click **Config** on the home page.
2. **IPSO supports both the IPv4 and IPv6 protocols. Select either the IPv4 or IPv6 protocols.**
 - a. For IPv4 ACLs, click the *Access List Configuration* link under the **Traffic Management** section.
 - b. For IPv6 ACLs, click the *IPv6* link. This takes you to the IPv6 page. Click the *Access List Configuration* link under the **Traffic Management** section.
3. **Click the link for the appropriate access control list in the ACL Name field.**

Expected outcome

You are taken to the page for that access control list.

The following items can be modified:

- Action
- Aggregation Class
- Source IP Address
- Source Mask Length
- Destination IP Address
- Destination Mask Length
- Source Port Range

Note

You can specify the source port range only if the selected protocol is either "any," 6, TCP, 17, or UDP.

- Destination Port Range
-

Note

You can specify the destination port range only if the selected protocol is either "any," 6, TCP, 17, or UDP.

- Protocol
 - TCP-Establishment flag—When it is selected, traffic matches this rule when it is part of the initial TCP handshake. This option applies only to IPv4 ACLs.
-

Note

You can specify the TCP Establishment flag only if the selected protocol is TCP, 6, or "any."

- Type of Service (TOS) for IPv4; Traffic Class for IPv6
 - DiffServ codepoint (DSfield)
-

Note

RFC 791 states that the least significant two bits of the DiffServ codepoint are unused. Thus, the least significant two bits for any value of the DSfield that you enter in the ACL rule will be reset to 0. For example, if you enter 0xA3, it will be reset to 0xA0 and the corresponding packets will be marked as 0xA0 and not 0xA3.

- Logical Queue Specifier (QueueSpec)
-

Note

The DSfield and QueueSpec field can be configured only when the rule's action is set to "prioritize."

4. **To modify the aggregation class, go to Associating aggregation classes with rules.**
5. **Modify the values in one or more of the edit boxes or drop-down window or (de)select a radio button. Click Apply.**
6. **To make your changes permanent, click Save.**

Further information

See *Access control list rules* , *Adding new rules to access control lists* and *Deleting access control list rules*

9

Deleting access control list rules



Steps

1. Click **Config** on the home page.
2. **IPSO supports both the IPv4 and IPv6 protocols. Select either the IPv4 or IPv6 protocols.**
 - a. For IPv4 ACLs, click the *Access List Configuration* link under the **Traffic Management** section.
 - b. For IPv6 ACLs, click the *IPv6* link. This takes you to the IPv6 page. Click the *Access List Configuration* link under the **Traffic Management** section.
3. **Click the link for the appropriate Access control list in the ACL Name field.**

Expected outcome

You are taken to the page for that Access Control List.

4. **Click the Delete check box next to the rule that you want to delete. Click Apply.**
5. **To make your changes permanent, click Save.**

Further information

See *Access Control List rules configuration in Voyager, Adding new rules to access control lists* and *Modifying access control list rules*

10

Aggregation classes

An aggregation class (AGC) is used to determine whether the traffic stream meets certain throughput goals. Traffic that meets these goals is conformant. Traffic that does not meet these goals is non-conformant. Depending on the configuration of the classifier rules, non-conformant traffic may be delayed, policed; that is dropped, or marked. An aggregation class groups traffic from distinct rules and measures its throughput.

You can configure an aggregation class with two parameters: *meanrate* and *burstsize*. The *meanrate* is the rate, in kilobits per second (kbps), to which the traffic rate should be coerced when measured over a long interval. The *burstsize* is the maximum number of bytes that can be transmitted over a short interval.

When you initially create an AGC, a burst of traffic is conformant—regardless of how quickly it arrives—until the size of the burst (in bytes) is equal to or larger than the *burstsize* you configured for the AGC. When the burst reaches the configured *burstsize*, traffic is non-conformant, but the AGC increases the rate at which traffic is transmitted based on the configured *meanrate*. Traffic that arrives consistently at a rate less than or equal to the configured *meanrate* will always be marked conformant and will not be delayed or dropped in the respective shaper or policer stages.

For more information, see *Creating aggregation classes* and *Associating aggregation classes with rules*

11

Creating aggregation classes



Steps

1. **Click Config on the home page.**
2. **Go to the *Aggregation Class Configuration* page.**

Either click the *Aggregation Class Configuration* link under the **Traffic Management** section, or click the *IPv6* link and then click the *Aggregation Class Configuration* link under the **Traffic Management** section.

3. **Enter the name of the aggregation class in the Name edit box in the Create A New Aggregation Class section.**
4. **Enter the bandwidth in the Mean Rate (Kbps) edit box.**
5. **Enter the burstsize in the Burstsize (bytes) edit box.**
6. **Click Apply.**

Expected outcome

The aggregation class you have just created appears in the **Existing Aggregation Classes** section.

7. **To make your changes permanent, click Save.**

Further information

See *Aggregation classes*, *Associating aggregation classes with rules* and *Deleting aggregation classes*

12 Deleting aggregation classes



Steps

1. Click **Config** on the home page.
2. Go to the *Aggregation Class Configuration* page.

Either click the *Aggregation Class Configuration* link under the **Traffic Management** section, or click the *IPv6* link and then click the *Aggregation Class Configuration* link under the **Traffic Management** section.

3. Click the **Delete** check box next to the aggregation class that you want to delete. Click **Apply**.

Expected outcome

The aggregation class you have just created appears in the **Existing Aggregation Classes** section.

4. To make your changes permanent, click **Save**.

Further information

See *Aggregation classes*

13

Associating aggregation classes with rules



Steps

1. Click **Config** on the home page.
2. **IPSO supports both the IPv4 and IPv6 protocols. Select either the IPv4 or IPv6 protocols.**
 - a. For IPv4 ACLs, click the *Access List Configuration* link under the **Traffic Management** section.
 - b. For IPv6 ACLs, click the *IPv6* link. This takes you to the IPv6 page. Click the *Access List Configuration* link under the **Traffic Management** section.
3. **Click the link for the appropriate access control list in the ACL Name field.**

Expected outcome

You are taken to the page for that access control list.

4. **Select SHAPE or PRIORITIZE from the Action drop-down window. Click Apply.**
5. **Select an existing aggregation class from the Aggregation Class drop-down window. Click Apply.**

Note

If there is no aggregation class listed, you need to create an aggregation class. Go to *Creating aggregation classes*.

Note

A rule treats traffic as if it were configured for "skip," if the traffic matches a rule whose action has been set to "prioritize" or "shape" and no aggregation class is configured.

6. To make your changes permanent, click Save.

Further information

See *Aggregation classes*

14 Queue classes

Queue classes (QCs) are used to instantiate a framework, or template, for output queue schedulers. Like access control lists (ACLs) they are created and configured and then associated with an interface.

There are a maximum of 8 priority-level queues for a QC. You can configure the size (in packets) of each queue level as well as the queue specifier. The queue specifier is a tag assigned by the classifier and is used as a key to look up the proper queue level. IPSO supports two different queuing types: strict priority queuing (SPQ) and weighted round robin (WRR), which uses assured forwarding (AF) PHBs. See RFC 2597 assured forwarding PHB group for more information.

Certain queue levels are pre-defined. The remaining queue levels can be assigned any name and QueueSpec you want. The following table shows the pre-defined queue values for SPQ:

<i>Name of Queue Level</i>	<i>Priority</i>	<i>IETF recommended DiffServ Codepoint</i>	<i>Queue Specifier Value</i>
Internetwork control	0	0xc0	7
Expedited forwarding	1	0xb8	6
Best effort	7	0	0

The following table shows the pre-defined queue values for WRR:

<i>Name of Queue Level</i>	<i>Priority</i>	<i>IETF recommended DiffServ Codepoint</i>	<i>Queue Specifier Value</i>
Internetwork control	0	0xC0	7
Expedited forwarding	1	0xB8	6

Assured forwarding 41	2	0x88	5
Assured forwarding 31	3	0x68	4
Assured forwarding 21	4	0x48	3
Assured forwarding 11	5	0x28	2
Best effort	7	0	0

When you configure an ACL rule to use the priority action, you must *configure an aggregation class (AGC)* . This AGC will function as a policer, that is, non-conforming traffic will be dropped.

You should configure the AGCs so that the aggregate of the NC and EF flows consumes no more than 50% of the output link bandwidth. This action prevents lower-priority traffic from being starved. See RFC 2598 for more information. The other policers should also be configured to prevent the lower-priority queue from being starved.

Internetwork control traffic, such as routing messages and keepalives, should be configured to use the IC queue so that it receives precedence over regular IP traffic. Note that locally originated internetwork control traffic is automatically sent through this queue. See RFC 791 for more information about internetwork control traffic.

A queue class can be configured to maximize device throughput or to minimize prioritized traffic latency. The QoS functionality is not achieved without a cost. The choice of QoS with minimal latency is the most costly in terms of forwarding performance, but it allows the least amount of head-of-line blocking for high priority traffic.

For more information, see *Creating queue classes* and *Setting or modifying queue class configuration values*

15 Creating queue classes



Steps

1. Click **Config** on the home page.
2. Go to the *Queue Class Configuration* page.

Either click the *Queue Class Configuration* link under the **Traffic Management** section, or click the *IPv6* link and then click the *Queue Class Configuration* link under the **Traffic Management** section.

3. To create a new queue class, enter its name in the **Create A New Queue Class** edit box.

Expected outcome

The new queue class appears in the **Existing Queue Classes** field.

4. Click **Apply**.
5. To make your changes permanent, click **Save**.

Further information

See *Queue classes*, *Deleting queue classes* and *Setting or modifying queue class configuration values*

16 Deleting queue classes



Steps

1. Click **Config** on the home page.
2. Go to the *Queue Class Configuration* page.

Either click the *Queue Class Configuration* link under the **Traffic Management** section, or click the *IPv6* link and then click the *Queue Class Configuration* link under the **Traffic Management** section.

3. Click the **Delete** check box in the **Existing Queue Classes** field next to the name of the Queue class you want to delete.

Expected outcome

The queue class disappears from the **Existing Queue Classes** field.

4. Click **Apply**.
5. To make your changes permanent, click **Save**.

Further information

See *Queue classes*

17

Setting or modifying queue class configuration values



Steps

1. Click either **Config** on the Voyager home page or click the **Traffic Management** link on the home page.
2. Go to the **Queue Class Configuration** page.

Either click the *Queue Class Configuration* link under the **Traffic Management** section, or click the *IPv6* link and then click the *Queue Class Configuration* link under the **Traffic Management** section.

3. Enter a name for each queue you want to configure in the **Logical Name** edit box.

Expected outcome

The name appears on the queue monitoring page.

4. To modify an existing queue class, in the **Existing Queue Classes** field, click on the name of the queue class you want to edit.

Note

Choose a name (with no spaces) that will allow you to identify the queue's purpose.

Note

Each queue class can have up to eight queues.

Strict priority queuing (SPQ) has three queues reserved for internetwork control, expedited forwarding, and best effort traffic.

Weighted round robin (WRR) has six queues reserved for internetwork control, expedited forwarding, assured forwarding 41, assured forwarding 31, assured forwarding 21, assured forwarding 11 and best effort.

- 5. Each queue class scheduling principle defaults to the strict priority queue type (SPQ), and can be changed to weighted round robin (WRR) by using the TYPE box pull down menu. Click Apply.**
-

Note

With the strict priority queue scheduling principle, all fields (LOGICAL NAME, QUEUE SPECIFIER, and MAX QUEUE LENGTH) can be modified at queue class priority levels 2-6. For the queue class priority levels 0,1, and 7, the modifiable field is only MAX QUEUE LENGTH.

Note

With the weighted round robin queue scheduling principle, all fields (LOGICAL NAME, QUEUE SPECIFIER, MAX QUEUE LENGTH, and WEIGHT) can be modified at queue class priority level 6. For the queue class priority levels 0-5 and 7, the modifiable fields are only MAX QUEUE LENGTH and WEIGHT.

- 6. Enter an integer for the logical identifier used to address each queue you configure within a queue class in the Queue Specifier edit box.**
- 7. For each queue, enter a value for the maximum number of packets that can be queued before packets are dropped in the Max Queue Length edit box.**

A value of zero (0) is used to disable a queue. Neither the network control nor the best effort queue can be disabled.

- 8. Configure a weight for the queue levels**

With WRR you can configure different weights (0-8) for the queue levels.

The queue weights define the proportions in which the link capacity is to be divided amongst the different classes. For example, if one slot has the weight 1 and another slot has the weight 2, then the second slot will get twice as much as the first flow.

9. **Click Apply.**
 10. **To make your changes permanent, click Save.**
 11. **To change the name of a modifiable queue level, enter the new name in the Logical Name edit box.**
-

Note

Choose a name (with no spaces) that will allow you to identify the queue's purpose.

Expected outcome

This name appears in the queue monitoring page.

12. **Click Apply.**
13. **To make your changes permanent, click Save.**

Further information

See *Queue classes* and *Associating queue classes with interfaces*

18

Associating queue classes with interfaces



Steps

1. Click **Config** on the home page.
2. Go to the *Queue Class Configuration* page.

Either click the *Queue Class Configuration* link under the **Traffic Management** section, or click the *IPv6* link and then click the *Queue Class Configuration* link under the **Traffic Management** section.

3. To associate a queue class with an interface, click on the appropriate physical interface in the **List Of Available Physical Interfaces** field.

Expected outcome

You are taken to the physical interface page for the interface you selected.

4. To enable QoS queuing, select either **MAX THROUGHPUT** or **MIN QOS LATENCY** from the **Queue Mode** drop-down window in **Queue Configuration** field.
5. Click **Apply**
6. Select the configured queue class you want to associate with the interface from the **Queue Class** drop-down window in the **Queue Configuration** field.

Note

If you do not select a queue class, the default class will be used. The default queue class has two queues, internetwork control and best effort.

7. **Click Apply.**
8. **To make your changes permanent, click Save.**

Further information

See *Queue classes* and *Setting or modifying queue class configuration values*

19 ATM QoS

ATM networks can provide different quality of service for network applications with different requirements. Unspecified bit rate (UBR) service does not make any traffic related guarantees. It does not make any commitment regarding cell loss rate or cell transfer delay. Constant bit rate (CBR) service provides continuously available bandwidth with guaranteed QoS.

The implementation supports CBR channels through a mechanism on an ATM network interface card (NIC) that limits the cell rate for each virtual channel you configure. The CBR feature limits the peak cell rate for each CBR channel in the output direction only. Each ATM port supports up to 100 CBR channels with 64 kbits/sec of bandwidth resolution.

20 Creating QoS descriptors



Steps

1. **Click Config on the home page.**
2. **Click the *ATM QoS Descriptor Configuration* link in the Traffic Management section.**
3. **To create an ATM QoS Descriptor, enter its name in the Create A New ATM QOS Descriptor edit box.**

The category for any new ATM QoS descriptor that you configure is set to constant bit rate (CBR).

CBR limits the maximum cell output rate to adhere to the requirements on CBR traffic imposed by the network.

Note

The default ATM QoS descriptor is set to unspecified bit rate; this descriptor cannot be modified.

-
4. **Enter a value for the maximum cell rate to be used in the output direction on a CBR channel in the Peak Cell Rate edit box.**

The peak cell rate is rounded down to a multiple of 64 kilobits/sec. One cell per second corresponds to 424 bits/sec.

Note

You can configure no more than 100 CBR channels per interface. The sum of the peak cell rate of all the CBR channels on an interface cannot exceed 146Mbs.

5. **Click Apply.**

Expected outcome

The new ATM QoS descriptor appears in the **Existing ATM QoS descriptors** field.

6. **Click Save to make your changes permanent.**

Further information

See *ATM QoS, Associating an ATM QoS descriptor with an interface and a virtual channel* and *Deleting ATM QoS descriptors*

21 Deleting ATM QoS descriptors

Before you start

Note

You can delete an existing ATM QoS descriptor only after you dissociate it from an existing permanent virtual channel (PVC).



Steps

1. **Click Config on the home page.**
2. *If the ATM QoS descriptor that you want to delete is associated with an existing PVC*

Then

complete the following steps:

- a. Click *Interfaces* link.
 - b. Click the appropriate ATM interface link in the **Physical** field.
You are taken to the physical interface page for the interface you selected.
 - c. Click the *ATM QoS Configuration* link. You are taken to the ATM QoS Configuration page for the physical interface you selected.
 - d. In the **QOS Configured PVCS** field, click the **QOS Descriptor** drop-down window and select **Default (UBR)**.
 - e. Click **Apply**, and then click **Save** to make your changes permanent.
3. **Click the *ATM QoS Descriptor Configuration* link in the Traffic Management section.**
 4. **In the Existing ATM QOS Descriptors field, click the Delete check box next to the name of the ATM QoS Descriptor that you want to delete.**

5. Click Apply.**Expected outcome**

The ATM QoS descriptor disappears from the **Existing QOS Descriptors** field.

6. Click Save to make your changes permanent.**Further information**

See *ATM QoS*

22

Associating an ATM QoS descriptor with an interface and a virtual channel

Before you start

Note

You cannot delete or modify a QoS descriptor that has been associated with a permanent virtual channel (PVC). You must first disassociate the PVC from the QoS descriptor. See *Deleting ATM QoS descriptors* for more information.



Steps

1. **Click Config on the home page.**
2. **Click *Interfaces* link.**
3. **To associate an ATM QoS descriptor with an interface, click the appropriate interface link in Physical field.**

Expected outcome

You are taken to the physical interface page for the interface you selected.

4. **Click the *ATM QoS Configuration* link.**

Expected outcome

You are taken to the ATM QoS Configuration page for the physical interface you selected.

5. **In the *Configure A New PVC* field, enter the virtual path identifier/virtual channel identifier (VPI/VCI) of the permanent virtual channel (PVC) you want to configure, in the VPI/VCI edit box.**

- 6. In the Configure A New PVC field, click the QoS Descriptor drop-down window and select the QoS descriptor with which you want to associate the PVC you configured.**
-

Note

You can change the QoS configuration of a PVC while it is being used. However, doing so results in a short break in traffic because the PVC is closed while QoS configuration values change. Afterward, the system reopens the PVC.

- 7. Click Apply.**

Expected outcome

The name of the new PVC and ATM QoS descriptor with which you associated the PVC appear in **QoS Configured PVCS** field.

- 8. Click Save to make your changes permanent.**

Further information

See *ATM QoS*

23 Common open policy server (COPS)

The common open policy server (COPS) provides a standard for exchanging policy information in order to support dynamic quality of service (QoS) in an IP (Internet protocol) network. This information is exchanged between PDPs (policy decision points) and PEPs (policy enforcement points). The PDPs are network-based servers that decide which types of traffic (such as voice or video) receive priority treatment. The PEPs are routers that implement the decisions made by the PDPs. In the Nokia implementation, the Nokia platform functions as a PEP.

For more information, see *Configuring COPS client IDs and policy decision points* and *Assigning roles to specific interfaces*

24

Configuring COPS client IDs and policy decision points

Purpose

You must configure at least one COPS client ID and a corresponding policy decision point, that is, policy server, for the COPS policy module to function.



Steps

1. Click **Config** on the **Voyager** home page.
2. Click the **COPS** link in the **Traffic Management** section.
3. In the **Configured COPS Modules** section click the **Diffserv PIB** link.

Expected outcome

You are taken to the COPS Diffserv specific configuration page.

4. In **Diffserv PIB Specific Configuration** section, enter the name of the new client ID in the **Create A New Client ID** edit box.

Note

You can configure multiple client IDs. Only one client ID can be active at a time.

5. Click **Apply**.

Expected outcome

The name of the new COPS client appears in a client ID list in the **COPS Security Configuration** section.

6. **To select an active client ID, click on the Client ID drop-down window and select a client name.**
7. **Click Apply.**
8. **Enter either the IP address or domain name the server to act as the policy decision point (PDP) in the Primary PDP edit box.**
9. **(Optional) Enter the IP address or domain name of the server to act as the secondary policy decision point (PDP) in the Secondary PDP edit box. Click Apply.**
10. **To make your changes permanent, click Save.**

Further information

See *Common open policy server (COPS)*, *Configuring security parameters for COPS client IDs*, *Activating and deactivating the COPS client*, *Changing client IDs associated with specific Diffserv configurations* and *Deleting client IDs*

25

Configuring security parameters for COPS client IDs

Purpose

The Nokia implementation lets you configure send and receive key IDs for each COPS Client ID to authenticate sessions with the PDP, or policy server.



Steps

1. Click either **Config** on the Voyager home page or click the **Traffic Management** link on the home page.
2. Click the **COPS** link in the **Traffic Management** section.
3. In the **Configured COPS Modules** section click the **Diffserv PIB** link.

Expected outcome

You are taken to the COPS Diffserv specific configuration page.

4. In the **COPS Security Configuration** section, click on the link for the name of the COPS client ID for which you want to configure security.

Expected outcome

You are taken to the COPS Security Configuration page for that client.

5. In the **SEQUENCE NUMBER** edit box, enter a value between 1 and 2147483647 to define the sequence number used for the COPS protocol.
6. Click **Apply**.
7. In the **Key ID** field, enter a value between 1 and 2147483647 in the **Send** edit box to define the send key ID used for the COPS protocol.

8. In the **Key** field, enter a string value of up to 64 characters in the edit box next to the **Send Key ID** value. This value defines the key used for the COPS protocol. Use alphanumeric characters only. Click **Apply**.
 9. In **Key ID** field, enter a value between 1 and 2147483647 in the **RECV** edit box to define the receive key ID used for the COPS protocol.
 10. In the **Key** field, enter a string value of up to 64 characters in the edit box next to the **RECV Key ID** value. This value defines the key used for the COPS protocol. Use alphanumeric characters only.
-

Note

You can configure up to 5 receive key IDs.

11. Click **Apply**.
12. To make your changes permanent, click **Save**.

Further information

See *Common open policy server (COPS)* and *Activating and deactivating the COPS client*

26 Assigning roles to specific interfaces

Purpose

The Nokia COPS implementation lets you assign roles to specific interfaces. A role refers to a logical name assigned to a group of objects within a network. The role name lets you group objects to which you want to assign a particular policy. You can also assign a combination of roles to a particular logical interface. You then apply policies to role(s) and not just to a single object.



Steps

1. **Click either *Config* on the Voyager home page or click the *Traffic Management* link on the home page.**
2. **Click the *COPS* link in the *Traffic Management* section.**
3. **In the *Interface Role Combinations* section, enter the name for a role in the edit box next to the appropriate logical interface name.**

The role name can be up to 31 characters long. Use alphanumeric characters, the period, hyphen or underscore symbols only. Do not begin a role name with the underscore symbol.

Note

You can assign multiple roles to each interface.

Note

You can assign different roles to different interfaces on the same system.

4. **Click *Apply***

5. To make your changes permanent, click Save.

Further information

See *Common open policy server (COPS)*

27

Activating and deactivating the COPS client

Purpose

You must activate the COPS client to implement the COPS module you configure.

You can deactivate the COPS client to halt the COPS module implementation.

You can maintain any existing module and role configuration. This configuration remains available if you reactivate the COPS client.



Steps

1. **Click either *Config* on the Voyager home page or click the *Traffic Management* link on the home page.**
2. **Click the *COPS* link in the *Traffic Management* section.**
3. **To activate or deactivate the COPS client, click the *Start* or *Stop* button in the COPS Client field.**
4. **Click *Apply*.**
5. **To make your changes permanent, click *Save*.**

Further information

See *Common open policy server (COPS)*, *Configuring COPS client IDs and policy decision points*, *Configuring security parameters for COPS client IDs*, *Changing client IDs associated with specific Diffserv configurations* and *Deleting client IDs*

28

Changing client IDs associated with specific Diffserv configurations

Purpose

You can change a client ID on a running system. Typically, each client ID refers to a specific policy or set of policies.



Steps

1. Click either **Config** on the Voyager home page or click the **Traffic Management** link on the home page.
2. Click the **COPS** link in the **Traffic Management** section.
3. Click the **Diffserv PIB** link in the **Configured COPS Module** section.

Expected outcome

You are taken to the COPS Diffserv specific configuration page.

4. In the **Diffserv PIB Specific Configuration** section, click the **Client ID** drop-down window and select the client ID name you now want to run.

Note

A list of all existing Client IDs appears in the **COPS Security Configuration** section.

Expected outcome

The name of the client ID you selected now appears in the **Client ID** field.

5. Click **Apply**.

6. To make your changes permanent, click Save.

Further information

See Common open policy server (COPS), Configuring COPS client IDs and policy decision points and Deleting client IDs

29 Deleting client IDs

Before you start

You cannot delete an active client ID. Deactivate the client ID before deleting it.



Steps

1. Click either **Config** on the Voyager home page or click the **Traffic Management** link on the home page.
2. Click the **COPS** link in the **Traffic Management** section.
3. Click the **Diffserv PIB** link in the **Configured COPS Module** section.

Expected outcome

You are taken to the COPS Diffserv specific configuration page.

4. Click the **Client ID** drop-down window in the **Diffserv PIB Specific Configuration** section and select either another existing client ID name or **NONE**.
5. Click **Apply**.

Expected outcome

You can now delete the client ID you disabled.

6. In the **COPS Security Configuration** section, click the **Delete** check box next to the name of the client ID you want to delete.
7. Click **Apply**.
8. To make your changes permanent, click **Save**.

Further information

See *Common open policy server (COPS)*

30 Example of rate shaping configuration

Purpose

The following example shows you how to limit ftp data traffic to 100 kilobits per second (kbps) with a 5000 byte burstsize on output interface eth-s2p1c0.



Steps

1. **Create an access control list.**
 - a. Click on **Config** in the home page.
 - b. Click on the *Access List Configuration* link under the **Traffic Management** section.
 - c. To create the access control list, enter its name in the **Create A New Access List** edit box, and click **Apply**.
 - d. Click the **Add Rule Before** check box next to the last rule, and click **Apply**.
 - e. Enter tcp in the **Protocol** edit box and enter **20** in both the **Source** or **Destination Port Range** edit box, and click **Apply**.
 - f. Select SHAPE from the **ACTION** drop-down window, and click **Apply**.

2. **Create an aggregation class.**
 - a. Click on the *Aggregation Class Configuration* link on the *Access Control List Configuration* page.
 - b. Enter the name of the new Aggregation Class in the **Name** edit box in the **Create A New Aggregation Class** section.
 - c. Click **Apply**, and then click **Save** to make your change permanent.
 - d. Enter **100** in the **Meanrate (KBPS)** edit box.
 - e. Enter **5000** in the **Burstsize (Bytes)** edit box.
 - f. Click **Apply**, and then click **Save** to make your change permanent.

3. **Associate the aggregation class with the rule you set when you created the access control list.**

- a. Click on the *Access Control List Configuration* link on the *Aggregation Class Configuration* page.
- b. For the rule you set up when you created the Access Control List, select the aggregation class you created from the **Aggregation Class** drop-down window.
- c. Click **Apply**.
- d. Select **ETH-S2P1C0** from the **Add Interfaces** drop-down window, and select **OUTPUT** from the **Direction** drop-down window.
- e. Click **Apply**, and then click **Save** to make your change permanent.

31 Example of expedited forwarding configuration

Purpose

This example illustrates the combined use of the access control list, traffic conditioning, and queuing features.

It demonstrates how to improve the response time to Telnet sessions between client and server systems over a private WAN connection within a corporate intranet as shown in the diagram below. The WAN interfaces for Network Application Platform (NAP) A and for Network Application Platform (NAP) B are ser-s3p1. The following configuration is done both on NAP A and NAP B.




Figure 1. Telnet session between client and server system over a private WAN connection within a corporate intranet



Steps

1. **Save the current configuration on each NAP before you set up QoS.**

This allows you to compare the relative performance of the QoS and non-QoS configurations.

- a. Click on **Config** in the home page.
 - b. Click on the *Manage Configuration Sets* link under the **System Configuration** section.
 - c. Enter pre-QoS in the **Save Current State To New Configuration Database** edit box.
 - d. Click **Apply**, and then click **Save** to make your change permanent.
- 2. Create an aggregation class**
- a. Click on **Config** in the home page.
 - b. Click on the *Aggregation Class Configuration* link under the **Traffic Management** section.
 - c. Enter **wan_1_ef** in the **Name** edit box in the **Create A New Aggregation Class** section.
 - d. Enter **100** in the **Meanrate (KBPS)** edit box.
 - e. Enter **5000** in the **Burstersize (Bytes)** edit box.
 - f. Click **Apply**, and then click **Save** to make your change permanent.
- 3. Create a queue class**
- a. Click on **Config** in the home page.
 - b. Click the *Queue Class Configuration* link under the **Traffic Management** section.
 - c. Enter **wan_1_ef** in the **Create A New Queue Class** edit box.
 - d. Click on the link to *wan_1_ef* in the **Existing Queue Classes** section to view existing queue class values.
-

Note

The queue specifier associated with expedited forwarding queue is 6.

- 4. Associate the wan_1_ef queue class with the appropriate interface.**
- a. Click on **Config** in the home page.
 - b. Click the *Interfaces* link.
 - c. Click on **SER-S3P1** in the **Physical** column.
 - d. In the **Queue Configuration** field, select MAX THROUGHPUT from the **Queue Mode** drop-down window, and click **Apply**.
 - e. In the **Queue Configuration** field, select **WAN_1_EF** from the **Queue Class** drop-down window.
 - f. Click **Apply**, and then click **Save** to make your change permanent.

5. **Create a new access control list rule to classify, condition, and prioritize Telnet traffic.**
 - a. Click on **Config** in the home page.
 - b. Click on the *Access List Configuration* link under the **Traffic Management** section.
 - c. Enter **wan_1_telnet** in the **Create A New Interfaces** edit box, and click **Apply**.
 - d. Select **SER-S3P1** from the **Add Interfaces** drop-down window.
 - e. Select **OUTPUT** from **Direction** drop-down window, and click **Apply**.
 - f. In the **Existing Rules FOR WAN_1_Telnet** section, click on the **Add New Rule Before** check box, and click **Apply**.
 - g. Select **PRIORITIZE** from the **Action** drop-down window, and then click **Apply**.
 - h. Select **WAN_1_EF** from the **Aggregation Class** drop-down window, and then click **Apply**.
 - i. For **NAP A**, enter **23** in the **Destination Port Range** edit box, and for **NAP B**, enter **23** in the **Source Port Range** edit box.
-

Note

The Telnet port number is 23.

- j. Enter **tcp** in the **Protocol** edit box; enter **0xB8** in the **DSField** edit box; and enter **6** in the **QUEUESPEC** edit box.
-

Note

0xB8 is the IETF differentiated-services codepoint (in hexadecimal) for expedited forwarding traffic.

- k. Click **Apply**, and then click **Save** to make your change permanent.

6. Test the configuration.

Start a Telnet session between the client and server.

7. Check the statistics on NAP A and NAP B

- a. Click **Config** on the home page.
- b. Click on the *Interfaces* link.
- c. Click on the link for *SER-S3P1* in the **Physical** column.
- d. Click on the *Interface Statistics* link.
- e. Scroll down to view statistics for Queue Class wan_1_ef.

Expected outcome

You should see values other than zero on both NAP A and NAP B for the **Packets Passed** and **Bytes Passed** counters in the **Expedited Forwarding** row.

8. **Use the Telnet session to generate traffic.**
9. **Check each NAP's interface statistics.**
 - a. Click **Config** on the home page.
 - b. Click on the *Interfaces* link.
 - c. Click on the link for *SER-S3P1* in the **Physical** column.
 - d. Click on the *Interface Statistics* link.
 - e. Examine the statistics for input and output traffic and compare them to the statistics for Expedited Forwarding traffic.
10. **Start an ftp session to create heavy (non-Telnet) background traffic over the WAN.**

Note

The Telnet session remains responsive. Use a text editor to examine a file.

11. **Save the QoS routing configuration (See Step 1), and restore the non-QoS configuration.**

Compare the difference in responsiveness when there is heavy WAN traffic both with and without QoS routing.