

# Solaris<sup>™</sup> Bandwidth Manager 1.6

---

*Bandwidth Management for IP Networks*



Sun Microsystems, Inc.  
901 San Antonio Road  
Palo Alto, CA 94303  
1 (800) 786.7638  
1.512.434.1511

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, JavaBeans, Solaris, Solstice SunNet Manager, and Solstice Enterprise Manager are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

**RESTRICTED RIGHTS:** Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

**DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.**

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, JavaBeans, Solaris, Solstice Enterprise Manager, et Solstice SunNet Manager sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Please  
Recycle



Adobe PostScript

# Contents

---

Introduction and Overview .....	1
How Bandwidth Management Works .....	1
Product Architecture .....	2
Interface with Accounting and Billing .....	2
Customizing and Extending the Product .....	2
Topologies and Examples .....	2
The Need for Guaranteed Bandwidth .....	3
Managing Network Capacity .....	4
Current Limits .....	4
Service-Level Agreements .....	4
Service-Driven Networks .....	5
Solaris Bandwidth Manager .....	6
Key Features .....	6
Key Benefits .....	7
Who Needs It? .....	7
How Bandwidth Management Works .....	9
Provisioning Rules .....	9
Traffic Classification and Scheduling .....	10
Bandwidth Borrowing .....	11
How to Calculate Bandwidth Needs .....	12

Product Architecture .....	13
Administration Tool .....	15
The Policy Agent .....	15
The Packet Classifier .....	15
URL Filtering .....	16
The Packet Scheduler .....	16
Type-of-Service Support .....	16
Monitoring, Accounting, and Billing .....	18
Statistics .....	19
Flows .....	19
Monitoring Statistics .....	19
Interface with Accounting and Billing .....	19
Customizing and Extending the Product .....	22
Installation Environments .....	23
Configuration Choices .....	23
High Availability (HA) .....	24
Integration with Directory Services .....	24
Typical Topologies .....	26
Case Study 1: Web Hosting .....	26
Case Study 2: Point of Presence .....	27
Case Study 3: Cross Domain .....	27
Configuration Planning Example .....	27
Additional Resources .....	30

## Introduction and Overview

---

Current networking technology has two major limitations:

- Availability of bandwidth cannot be predicted in terms of quantity or quality. Applications that require a specific quality of service, specifically bandwidth and delay, cannot currently reserve the quality of service they need.
- It is difficult to control which applications or users get a share of the available bandwidth because it is impossible to differentiate between different types of traffic. An application or a user can take control of all the available bandwidth, preventing other applications or users from using the network.

Solaris™ Bandwidth Manager software controls the bandwidth allocated to applications, users, and organizations sharing the same intranet or Internet link. It enables Internet service providers (ISPs) and MIS departments to:

- Gain full control of network traffic
- Ensure bandwidth for mission-critical business applications
- Reduce traffic congestion and increase network efficiency
- Control user and application access to network resources
- Gather detailed network-use statistics and accounting data to enable billing for resources used
- Provide differentiated classes of service to users

## How Bandwidth Management Works

Solaris Bandwidth Manager works by sorting traffic into classes based on the application type, source and destination addresses, URL, and type-of-service field. It then schedules the traffic according to the minimum and maximum bandwidth defined for each class.

# Product Architecture

Solaris Bandwidth Manager contains these major components:

- The *administration tool*, which provides a graphical interface for configuring bandwidth allocation
- The *policy agent*, which implements the configuration and handles communication with the kernel module
- The *kernel module*, which is viewed as a STREAMS module by the IP stack and contains the packet classifier and packet scheduler
- The Java™ technology, C statistics, and accounting *APIs*
- The *configuration API*

## Interface with Accounting and Billing

Statistics and accounting data extracted from Solaris Bandwidth Manager can be input into billing systems. Four output possibilities exist: ASCII format, NetFlow, Java and C APIs, or SNMP.

Accounting data includes details such as sender and receiver IP addresses, packet and byte counts, time stamps, and application port numbers. Thus, both Internet service providers and enterprise customers may choose from a variety of charging parameters including time of day, bandwidth usage, application usage, traffic class, and so on.

## Customizing and Extending the Product

The Java and C APIs included with Solaris Bandwidth Manager enable the development of JavaBeans™ plug-ins or applications. The JavaBeans plug-ins are placed directly in the Solaris Bandwidth Manager policy agent. The applications either access Solaris Bandwidth Manager remotely, or exist on the same local machine.

## Topologies and Examples

In addition to an overview of Solaris Bandwidth Manager, this white paper also provides examples of typical topologies and configuration planning examples.

# The Need for Guaranteed Bandwidth

---

Corporate networks are rapidly evolving from a classic client/server paradigm toward an intranet-based model, founded on information-sharing and Web navigation. The resulting increased use of the network is causing increased demand for bandwidth. Simply shackling on new connections and fatter pipes is not cost-effective, and won't guarantee availability where it's most needed.

The emerging intranet-based computing model implies many new factors. And MIS managers and service providers will have to take these into account when planning networks and offering services to their customers:

- Very different types of traffic coexist, with different requirements in terms of quality of service. For example, interactive telnet traffic, real-time audio and video traffic, and bulk FTP traffic.
- It is becoming increasingly common for mission-critical applications to be deployed on the intranet. These include enterprise resource planning (ERP) applications such as SAP and Oracle. Also, use of e-commerce and business-to-business transaction packages on extranets is increasing.
- On servers, key resources such as CPU and memory are managed by system management tools. Network bandwidth, which is an even more essential resource, also requires management.
- There is demand for increased amounts of bandwidth: People may stay on the link for extended periods of time and download large amounts of data.
- There is demand for guaranteed quality of service in terms of bandwidth and minimum delay. Emerging Internet applications are both bandwidth-intensive and time-sensitive, often requiring support for voice, video, and multimedia applications across the network infrastructure.
- User attitudes are changing: Users expect instant access to information without delays or restrictions, especially if that information is critical to their work.

# Managing Network Capacity

Almost all network links are used by more than one user or application. This means that the available bandwidth must be shared between them. Bandwidth management tools provide management and control for how this is done.

If a network link is continuously congested, the link needs to be upgraded to provide greater capacity. In many cases, however, the typical load on a link is within the link capacity, and the link is congested only temporarily. Temporary congestion is sometimes predictable. For example, typically there are peaks in network use at certain times of the day or following an event. Other causes of temporary congestion, such as the transfer of a large file, are not possible to predict.

If the average use of a link is within the link capacity, considerable improvements in the performance of the network link can be made by managing how the available bandwidth capacity is used. Allocating bandwidth to a specific type of traffic helps to optimize the usage of the available bandwidth.

## Current Limits

Today, two major limitations exist in current networking technology:

- The bandwidth available on a link at any given moment cannot be predicted, in terms of quantity or quality. Bandwidth management is needed to allow applications that require a specific quality of service, namely bandwidth and delay, to reserve the bandwidth quality of service they need.
- It is difficult to control which applications or users get a share of the available bandwidth. In some circumstances, an application or a user can take control of all the available bandwidth, preventing other applications or users from utilizing the network. To solve this problem, companies can either add extra capacity at extra cost, resulting in an over-provisioned network that still doesn't guarantee equal access, or they can introduce bandwidth allocation.

## Service-Level Agreements

Virtual private networks (VPNs) are popular value-added Internet services, which corporations are increasingly utilizing. Enterprise customers seeking a VPN provider are more likely to sign with an ISP that can offer a contractual service-level agreement that guarantees quality of service (QoS).

While service-level agreements cannot yet guarantee end-to-end service across the public Internet, they can be implemented for transport over a single vendor network, or for Internet server hosting. In these areas, a service-level agreement is an important differentiator for an ISP.



Generally, customers subscribe to a particular class of service (for example, premium or standard), and sign a service-level agreement accordingly. Packet throughput is monitored as part of the agreement.

Value-added services, including security and hosting, are predicted to grow rapidly. ISPs that want a piece of the pie clearly need to implement bandwidth management in order to meet service-level agreements that guarantee QoS. Only efficient bandwidth management will enable them to tune network behavior so that users have the QoS they paid for.

As service-level agreements become standard, they will be required within corporations, as well. Departments will demand service-level agreements with the corporate MIS department.

## Service-Driven Networks

The new network service paradigm relies on what Sun Microsystems calls the service-driven network. This is a responsive, reliable, modular infrastructure based on the latest generation of management technology that is built on dynamic, flexible management services.

To respond to today's business needs, ISPs and large enterprises have no choice; they must deploy the service-driven network. It enables them to deliver innovative services (such as unified roaming, push browsers, multicast, and online shopping) to customers more quickly and at a lower cost than ever before.

Efficient bandwidth management is a crucial component of the service-driven network. When integrated with a distributed directory service and a policy-based approach to management, it provides a powerful, flexible solution for provisioning QoS across an entire network.

# Solaris Bandwidth Manager

---

Solaris Bandwidth Manager software controls the bandwidth allocated to particular applications, users, and organizations sharing the same intranet or Internet link. By allowing network administrators to guarantee minimum bandwidth and prioritize traffic, it enables them to prevent a small number of applications and users from consuming all available bandwidth and causing major network congestion.

Solaris Bandwidth Manager meets the needs of ISPs and MIS departments that need to offer their customers service-level agreements on a contractual basis, by enabling them to:

- Guarantee bandwidth for mission-critical business applications
- Reduce traffic congestion and increase network efficiency
- Control user and application access to network resources
- Gather detailed network-use statistics and accounting data which enables billing for resources used
- Provide differentiated classes of service to users

## Key Features

- Incoming and outgoing traffic is managed based on traffic type (telnet, FTP, e-mail, NFS), end-user source or destination address, and organization source or destination address.
- Bandwidth management rules can be configured to map organization, systems, or geographical layouts.
- Can be integrated with an LDAP directory server. User profiles and bandwidth management policies can be stored in the directory and dynamically retrieved by Solaris Bandwidth Manager.
- Any type of UDP/TCP-based traffic can be managed by Solaris Bandwidth Manager software, which works within a heterogeneous environment without any modification of the systems accessing the gateway.

- Runs on top of WAN links such as PPP, as well as LAN links such as Ethernet and FDDI.
- Works seamlessly with routers, providing TOS field marking and NetFlow data exporting.
- Java technology-based graphical user interface makes it easy to specify bandwidth allocation policies, and enables remote monitoring and configuration.
- Provides a full set of reporting utilities to track how bandwidth is used, and how efficiently the traffic provisioning rules are working. Reports can interface with most spreadsheet software.
- Detailed flow- and class-based accounting information is provided through ASCII output, NetFlow, and SNMP. Java and C APIs enable interfacing with billing applications.
- Customizable and extensible, the software comes with Java configuration APIs and a dynamically configurable policy agent based on the Java Dynamic Management™ Kit. Companies can develop Java plug-ins of their own, and slot these into the policy agent.

## Key Benefits

- Gives full control of network traffic to the network administrator, and prevents a small number of applications or users from consuming all the available bandwidth.
- Enables service providers to ensure adequate levels of service without over provisioning their network or overbooking bandwidth. As a result, it helps network administrators to get the most out of their existing network resources.
- Reduces the risk of network congestion.
- QoS policies can be defined corporate-wide, and dynamically adjusted according to local conditions.
- Integrates into existing, heterogeneous environments with no software or hardware upgrades required.

## Who Needs It?

### *LAN Administrators*

A LAN administrator leasing a network connection from a service provider can use Solaris Bandwidth Manager to make the most efficient use of the leased capacity.

Bandwidth management makes sure that the higher-priority traffic is sent first, and also helps ensure the maximum use of the capacity paid for. It is no longer necessary to over-specify requirements just to guarantee that priority traffic can be sent. In fact, LAN administrators may even be able to reduce the leased capacity.

### *MIS Managers*

Solaris Bandwidth Manager enables MIS managers of large corporations to regulate the traffic in their network, providing network services to numerous clients.

A guaranteed minimum of bandwidth can be provided to a department, and as a bonus, additional bandwidth can be delivered when it is not required by other departments. This makes capacity planning both easier and more accurate.

### *Web Service Providers*

Solaris Bandwidth Manager enables Web service providers to guarantee client companies that a given bandwidth is available to customers visiting their Web sites.

Today, many Web providers' charges are based either on disk space usage or the number of times a site is visited. Disk space used is not a good indication of the cost to the provider, since a small site that is visited frequently can be as expensive to provide as a large site that is visited less frequently. Using the number of visits to a site is a better indicator of the cost to the provider, but is potentially an unbounded cost for the client. With Solaris Bandwidth Manager, Web service providers can charge clients for delivering guaranteed bandwidth to their sites.

### *Internet Service Providers*

Solaris Bandwidth Manager enables ISPs to provide various classes of services to customers. For example, they could offer Premium and Standard services, each with guaranteed minimum access levels, to suit the needs and budgets of their customers.

Solaris Bandwidth Manager extracts usage statistics through an accounting API. This information can then be used by billing package, to provide real-time, usage-based billing.

### *Enterprise Service Providers*

An enterprise service provider that is hosting applications and databases for commercial companies that want to outsource hosting of such mission-critical business tools can use Solaris Bandwidth Manager. It guarantees client companies that a given quality of service is available to users when they are accessing applications and databases. It also enables billing on a per-usage basis.

## How Bandwidth Management Works

---

Solaris Bandwidth Manager enables network administrators to manage the bandwidth used by IP traffic by:

- Allocating traffic to a class based on the application type, source and destination addresses, URL group, or a combination, then assigning individual limits for each class. For example, “traffic to engineering must have at least 50 percent of the link,” or “HTTP traffic cannot exceed 10 percent of the link.”
- Prioritizing traffic. Some types of traffic, for example, interactive traffic generated when using telnet or rlogin, need a quick response time. Solaris Bandwidth Manager can assign a higher priority to that traffic. Traffic that does not require a quick response time, such as a file transfer using FTP, can be assigned a lower priority.

By balancing the bandwidth allocated to different types of network traffic and the relative priorities, network performance can be optimized.

### Provisioning Rules

The provisioning rules used by Solaris Bandwidth Manager to classify traffic are based on the following filters:

- Traffic type, corresponding to TCP/IP or UDP ports or services (HTTP, FTP, e-mail, news, telnet, NFS, and so on)
- Source or destination IP address: Facilitates discrimination between traffic coming from different machines
- Network, subnetwork source, or destination IP address: Enables discrimination between specific organizations or based on network topology
- URL: Wild cards in the URL specification may be included
- Type-of-service (TOS) value

Filters can be combined using hierarchies. For instance, a specific QoS can be assigned to all the traffic originating from an organization. Within this traffic, a subset of the specified QoS is reserved for e-mail.

For the first four filters, groups of items can be specified, for example, a group of several URLs or several IP addresses.

### *Fine-Grained Filtering*

Bandwidth allocation based simply on filtering by protocol is not sufficient to meet bandwidth management needs.

One of the key issues in this area is the extensive use of HTML/HTTP systems for e-commerce. A fine level of granularity is needed for bandwidth management to take into account more than just the protocol when assessing the relative importance of network traffic. Bandwidth management must base allocation not only on protocol type, but also on the applications and users involved. Solaris Bandwidth Manager enables just this type of fine-grained filtering.


### *Provisioning Examples*

Some examples of the specifications which can be defined using provisioning include:

- Ensure that telnet traffic from workstation A, B, or C has a minimum allocated bandwidth of five percent over a leased line.
- Reserve 50 percent of a WAN link to traffic that originates from machines located in offices A, B, and C. Within this 50 percent, reserve 10 percent for NFS traffic.
- Ensure that FTP transfer from machine B in office Y will not use more than 30 percent of the capacity of the line between offices Y and X.

## Traffic Classification and Scheduling

Class definitions are hierarchical, and every class has a parent. For example, if a class for FTP traffic and a class for FTP traffic to a host called `pear` are defined, the classes are connected in a hierarchy as shown below.



In this example, the FTP class is a child of the root class and is the parent of the `ftp-to-pear` class.

The configuration of Solaris Bandwidth Manager software specifies the set of known classes for a network node defined in terms of some or all of these factors. It also allocates a percentage of bandwidth and a priority to each class. The priority of a class is an integer from 1 (highest priority) to 7 (lowest priority).

When a packet arrives, the packet classifier analyzes the packet protocol, source information, and destination information and allocates the packet to a class queue where it waits to be processed. If the queue to which a packets is allocated is full, the packet is dropped. Normal retransmission ensures that the packet is sent again.

The Solaris Bandwidth Manager scheduler uses the percentage bandwidth configured and the priority for each class to decide the order in which class queues are processed. Within a class queue, the packets are processed on a first-in, first-out basis. When the network traffic reaches the maximum allocated to a class, packets from the next class in priority order are processed.

## Bandwidth Borrowing

Each class is allocated a percentage of the available bandwidth, and when that limit is reached, normally no more traffic from that class can be forwarded. However, if the network link is not being fully used, a class can temporarily borrow bandwidth from its parent class, and send traffic at a percentage that exceeds its allocation (up to a predefined maximum bandwidth).

Similarly, if the bandwidth of a child class is not being fully used, the parent of that class can use it. The amount of spare bandwidth allocated to a class depends on the percentage of bandwidth that the class can borrow. This is the maximum bandwidth.

Spare bandwidth is allocated among the highest priority classes first until there is no more spare bandwidth, or until those classes all reach their maximum bandwidth. If this happens, then the remainder of the spare bandwidth is divided between the next highest priority classes, and so on.

It is possible to define a class that has zero-percent bandwidth bandwidth allocated but may borrow bandwidth from its parent class. A packet allocated to such a class is forwarded only if there is no other traffic of higher priority waiting. Allocating zero percent and no borrowing to a class means that the class is blocked.

# How to Calculate Bandwidth Needs

There are no precise methods for choosing an initial configuration of Solaris Bandwidth Manager. In practice, using an iterative approach is often the best way.

One method starts with a configuration that reflects the actual usage patterns — perhaps adjusted to solve any serious problems, monitor the performance over a period, and fine-tune the configuration.

Alternatively, another method starts by allocating equal bandwidth to all classes, and then running in stats mode to monitor how traffic is classified, before allocating real bandwidth percentages to classes.

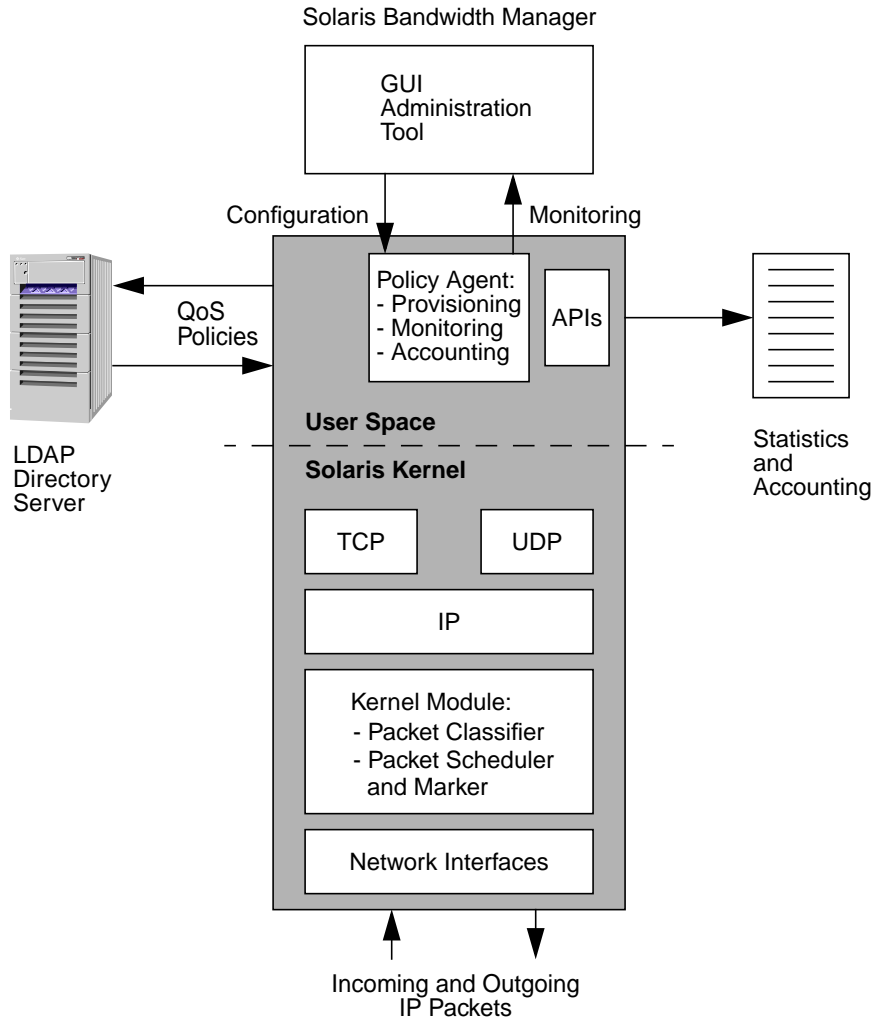


## Product Architecture

---

Solaris Bandwidth Manager contains the following major components:

- The *administration tool*, which provides a graphical interface for configuring bandwidth allocation. This can be run as an applet or application from any machine in the network that has a Java virtual machine.
- The *policy agent*, which implements the configuration and handles communication with the kernel module.
- The *kernel module*, which is viewed as a STREAMS module by the IP stack. It contains the classifier and the scheduler.
  - The *classifier* allocates packets to a class queue according to information in the IP or application header.
  - The *scheduler* sends the queued packets according to the predefined priority and transfer rate for each class. It can also mark the TOS fields of IP packets to a specific value.
- The *Java and C statistics and accounting APIs*, which provide an interface for gathering statistics and accounting information about packet classification and throughput. This data can be used as an input to a billing system.
- The *configuration API*, which enables modification of the Solaris Bandwidth Manager configuration from any custom application.



# Administration Tool

The administration tool can be used to configure Solaris Bandwidth Manager. It communicates with the kernel module through the policy agent, and has two modes of operation:

- In online mode, the configuration currently being used by the kernel module can be changed. This is useful if an immediate temporary change is required due to a problem in the network. The current configuration may also be saved to preserve any changes. With online mode, the consequences of a particular configuration can be observed before being saved.
- In offline mode, a configuration can be changed without disturbing the current behavior of the kernel module. This is useful for making changes in the configuration without disrupting users, and having the changes implemented the next time the product is restarted.

## The Policy Agent

The policy agent is the communications hub of Solaris Bandwidth Manager. It controls the information sent to and from all other components, as well as the policies that they operate.

## The Packet Classifier

The packet classifier collects packets from the IP layer. It applies the filters defined by the provisioning rules to assign each packet to a class. The class of a packet is defined by one or more of the following factors:

- IP source address
- IP destination address
- TCP or UDP source port
- TCP or UDP destination port
- Protocol (TCP or UDP)
- URL
- TOS value

This combination of factors enables very fine-grained filtering.

## URL Filtering

URL filtering uses wild cards to enable the following distinctions:

- Using \*.html, \*.gif, and \*.jpeg, the network administrator can distinguish between text-based information (\*.html) and images (\*.gif and \*.jpeg). This distinction enables prioritization of text loading and specification of which images will load in the background.
- Between different intranet pages, prioritizing crucial information such as price lists over archive files.
- Between different URLs to restrict access to nonwork-related sites. For example, the network administrator can prioritize access to `www.sun.*` over access to `www.games.*`

## The Packet Scheduler

The scheduler uses the percentage bandwidth configured and the priority for each class to decide the order in which class queues are processed. Within a class queue, the packets are processed on a first-in, first-out basis. When the network traffic reaches the maximum allocated to a class, packets from the next class in priority order are processed.

## Type-of-Service Support

An IP packet contains a type-of-service (TOS) field, which conveys information about how the packet should be processed. Solaris Bandwidth Manager can use this information when classifying a packet. It can also change the information to influence how the packet is processed.

The IP specification includes a definition of a TOS field in an IP packet header. This is intended to be used by upper-layer protocols to pass information to the Internet layer about how to optimize routing for the packet.

Network topology means that there are often a number of available routes between the source and destination of a packet. Some routes are more reliable than others. Some are expensive, while others are low cost but slow. The most suitable route for a packet depends on the application and user, and might even vary with other factors such as the time of day.

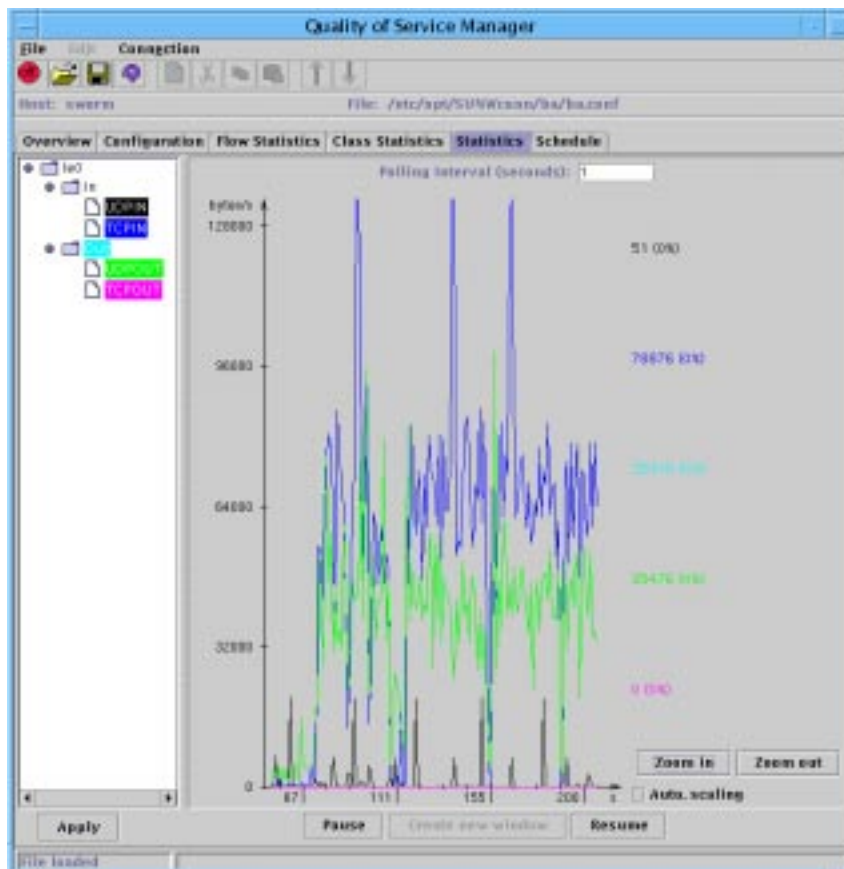
The Internet layer has no direct knowledge of how to optimize a route for a given application or user. The TOS facility is intended to provide hints about how best to route a packet, influencing both queueing algorithms and routing. The possible settings of the TOS field are:

- Minimize delay
- Maximize throughput
- Maximize reliability
- Minimize monetary cost
- Normal service

The TOS facility has not been widely used in the past, but the Internet Engineering Task Force (IETF) is now working to modify the definition of TOS and encourage its use. Solaris Bandwidth Manager can be configured to use the TOS field value as a classification criterion. The scheduler can override the TOS value if necessary when sending out a packet, using different priorities according to the predefined configuration.

# Monitoring, Accounting, and Billing

Solaris Bandwidth Manager provides real-time statistics on resource usage as well as service level agreement, commitment-level information (when combined with the right applications). This information can be used for monitoring network usage, or as input to accounting and billing systems.



## Statistics

Statistics information includes per-class and per-flow statistics on:

- Number of bytes and of packets sent
- Number of bytes and of packets dropped
- Average throughput
- Number of borrows from a class to its parent.

## Flows

A flow is a complete exchange of information between a sender and a recipient, as seen from the user's point of view. Examples of flows include sending a mail message or downloading a Web page.

Monitoring flows rather than classes gives a more accurate picture of network usage at finer granularity. This helps to predict future network needs more accurately, and provides information that can be used in accounting.

## Monitoring Statistics

There are several ways by which statistics and accounting information can be extracted from Solaris Bandwidth Manager:

- Statistics can be viewed directly and displayed graphically via the Java administration tool (for example, Sun Management Center). They can also be viewed in ASCII format via a command -line interface.
- Statistics may be viewed via any billing or accounting package compatible with version 5 of the CISCO NetFlow protocol.
- Java and C APIs can be used to build monitoring utilities or integrate Solaris Bandwidth Manager statistics into monitoring systems.
- A simple network management protocol (SNMP) management information base (MIB) is included with Solaris Bandwidth Manager. The MIB data can be extracted, processed, and viewed using an appropriate tool.

## Interface with Accounting and Billing

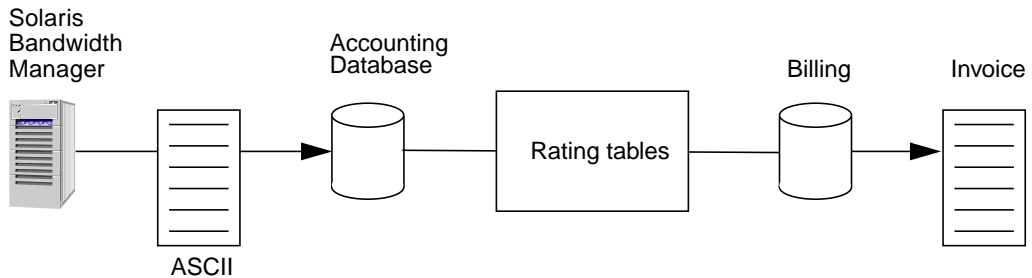
Statistics extracted from Solaris Bandwidth Manager can be input into the accounting function in billing systems. Four output possibilities exist:

- ASCII format

- NetFlow
- Java and C APIs
- SNMP

### *ASCII Format*

An accounting agent collects accounting information and outputs it in ASCII format, which can be accepted by a billing system.



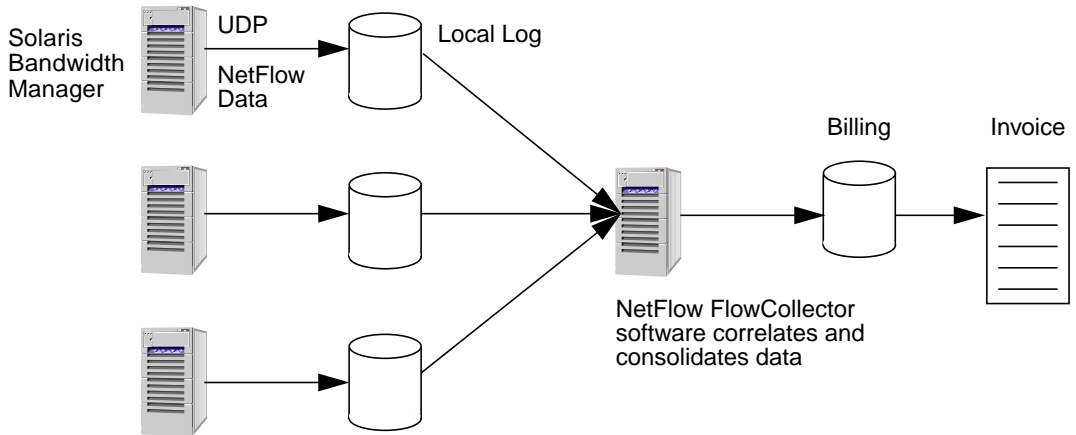
Accounting data includes details such as sender and receiver IP addresses, packet and byte counts, time stamps, and application port numbers. Thus, both Internet service providers and enterprise customers may choose from a variety of charging parameters, including time of day, bandwidth usage, application usage, and traffic class. Accounting data also provides the base for billing for advanced, QoS-enabled IP services.

### *NetFlow*

NetFlow support is built into Solaris Bandwidth Manager. NetFlow technology, developed by Cisco Systems, allows the collecting and processing of information for accounting and billing, network planning, and network monitoring.



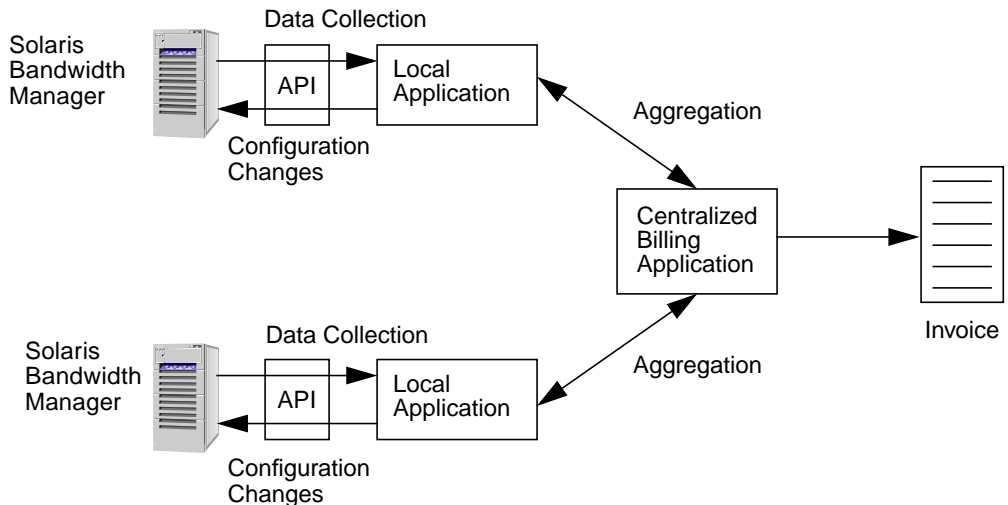
Solaris Bandwidth Manager can be configured to send NetFlow v5 datagrams directly to Cisco NetFlow FlowCollector software (which is supported by the Solaris Operating Environment).



### Java and C APIs

The inclusion of Java and C APIs with Solaris Bandwidth Manager allows for development of custom applications for real-time, usage-based billing.

For example, proprietary applications can be built for local collection of accounting data and aggregation into a centralized billing application.



The local application collects local accounting information and events (such as new flow and end of flow). It reacts to traffic thresholds and events, managing these locally or escalating events to the central billing application.

### *SNMP*

Traffic statistics and configuration information can be read by any standard network management application, such as Solstice SunNet Manager™ or Solstice Enterprise Manager™, via the SNMP interface included in Solaris Bandwidth Manager.

## Customizing and Extending the Product

The Java and C APIs included with Solaris Bandwidth Manager enable the development of JavaBeans plug-ins or applications.

The JavaBeans plug-ins are placed directly in the Solaris Bandwidth Manager policy agent. The applications either access Solaris Bandwidth Manager remotely, or exist on the same local machine.

### *Plug-ins*

Java plug-ins can be developed to implement network usage policies. For instance, a plug-in could perform internal statistics monitoring and react to statistics events by instigating an action. The action could be to generate an alert, or to dynamically change the configuration of Solaris Bandwidth Manager.

An example is a plug-in that checks traffic statistics. If it detects that Premium Class traffic is reaching capacity, it can dynamically reconfigure Solaris Bandwidth Manager. As a result, Standard Class traffic is disconnected to ensure guaranteed access to Premium Class traffic.

### *Applications*

Solaris Bandwidth Manager APIs can be used to develop management applications, or to interface with existing management applications. For example, an application could be developed that provides a graphical user interface to integrate Solaris Bandwidth Manager into a proprietary network management solution.

# Installation Environments

---


## Configuration Choices

Solaris Bandwidth Manager can be installed in two modes: Traffic Manager or Application Controller.

### *IP Transparent Mode: Traffic Management*


Solaris Bandwidth Manager can be installed in “IP-transparent” mode on a host between a LAN and a router. In this mode, it acts as a traffic manager, controlling access to the network (LAN, WAN, or Internet).

This mode is called IP-transparent because the host running Solaris Bandwidth Manager is completely transparent to the IP network, and is perceived as just another machine connected to the LAN. The LAN and the WAN behave as though they are directly connected through the router only. It is not necessary to modify the routing tables.



## *Server Mode: Application Control*

Installed on a host that is a source of IP traffic, Solaris Bandwidth Manager acts as an application controller, controlling output from a server to the network (LAN, WAN, or Internet) and corresponding incoming traffic. The server may be a file, application, or Web server. Traffic is controlled by server or by application.



## High Availability (HA)

When Solaris Bandwidth Manager is used in transparent mode, any failure of the hardware or software could result in link unavailability. Several standard solutions guarantee higher availability:

- Path redundancy with intelligent routing algorithms such as OSPF
- Router redundancy with Hot Standby Routing Protocol combined with a routing protocol that converges very rapidly (such as EIGRP)
- Link redundancy with link monitoring by fault-tolerant transceivers or switches, and use of spanning tree protocol

## Integration with Directory Services

Service level agreements, QoS guarantees, and other such services are becoming increasingly critical to the success of service providers in attracting and retaining clients. These services must generally be linked to the user rather than to a particular device. To provide guaranteed QoS on a user-by-user basis, information about users, network devices, and services has to be stored in a single, distributed data store.

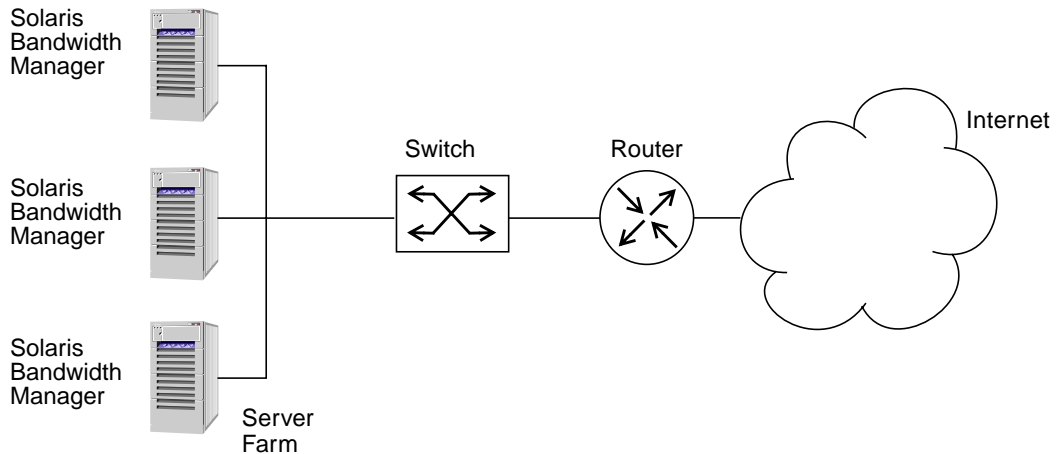
Solaris Bandwidth Manager configuration information and policy information can be stored using an LDAP directory service. As a result, a company must maintain only one directory of user profiles, which is accessible from anywhere on the network and by each instance of Solaris Bandwidth Manager. The benefits are obvious: reduced data administration overheads, greater accessibility of data, and improved accuracy.

# Typical Topologies

---


## Case Study 1: Web Hosting

This example shows Solaris Bandwidth Manager in a Web hosting environment. Traffic outgoing from the server farm is prioritized by server or application. This enables the ISP to guarantee bandwidth to each customer according to the terms of the hosting contract.




## Case Study 2: Point of Presence

In this example, Solaris Bandwidth Manager is being used to regulate the traffic generated by LAN-based PCs accessing the Internet.



## Case Study 3: Cross Domain

This example demonstrates Solaris Bandwidth Manager used to regulate use of the network to ensure that the branch office can access important applications and data on the corporate network. For example, this ensures that access to the corporate sales database is given priority over Internet access during core office hours.



## Configuration Planning Example

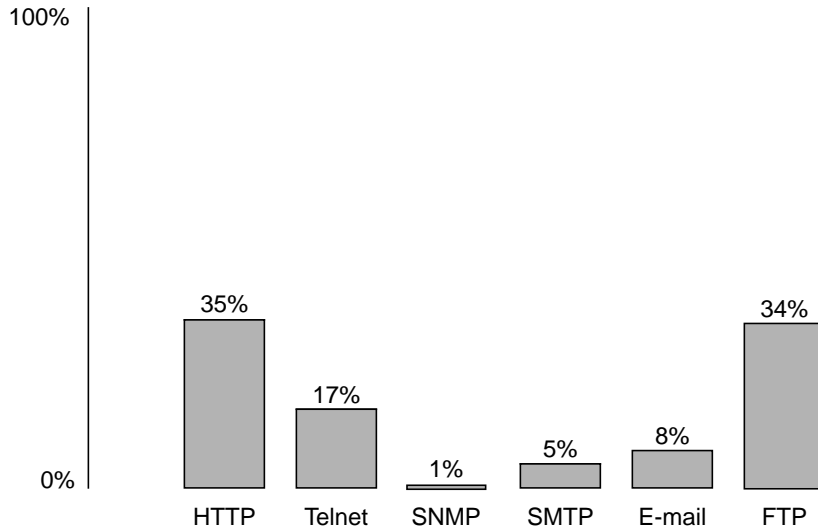
This section contains an example of bandwidth management being used within the European network of the Example Corporation at three points: Paris, Bonn, and London.

In the example, the Paris and Bonn sites each have a busy LAN, and route traffic from the LAN and from other sites on to the London site. From Paris to London there is a 256-Kbps line. From Bonn to London there is a 768-Kbps line. There is also a dial-up link directly from Paris to Bonn. London has its own LAN, and routes traffic from it, Paris, and Bonn to a U.S. site over a 10-Mbps line.

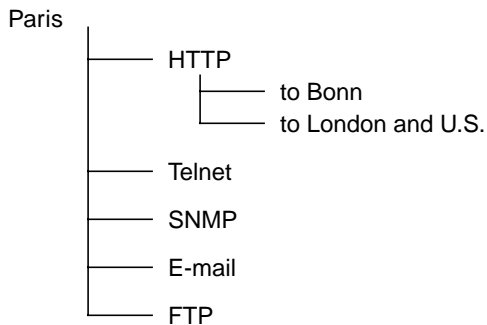
At all three sites, the network administrator monitored the actual network usage over a period of time, and asked users what they thought were the three most important uses of the network. The data for the Paris site is shown here, with the configuration that was designed.

### *Actual Network Use for Paris Site*

The network users at the Paris site consider e-mail, file transfer, and Web access to be the most important uses of the network. The actual usage pattern is shown below.



Using the data on network use and user input, the network administrator designed the class hierarchy shown below:





## *Bandwidth Allocation on the Paris Server*

<b>Class Description</b>	<b>Percentage Bandwidth Allocated</b>	<b>Priority</b>
Root	5	1
HTTP to Bonn	5	3
HTTP to London or U.S.	10	3
HTTP (to elsewhere)	5	5
Telnet	30	1
SNMP	5	1
E-mail	20	4
FTP	15	7
Default	5	7

The priorities and percentages of assigned bandwidth are shown in the table above.

This configuration is an example of how the allocated bandwidth and class priority can be combined to optimize network usage.

Bandwidth is allocated according to observation of current usage patterns, and constraints that have been decided upon, in order to improve service. The highest priority is given to traffic types that need access to available bandwidth when necessary. Usually, this is the case for interactive traffic and system monitoring traffic.

In this example, telnet and SNMP traffic are given the highest priority and a sufficient percentage of the guaranteed bandwidth. Note that if unused, this bandwidth can be borrowed by other traffic in priority order.

With this configuration, the bandwidth used by FTP traffic is constrained to 15 percent, contrasting with the current usage figure of over 30 percent. FTP is given the lowest priority.

## Additional Resources

---

For more information on how to configure and maintain Solaris Bandwidth Manager software in a wide variety of configurations and application environments, Sun provides:

- Solaris Bandwidth Manager Web Site ([www.sun.com/software/bandwidth](http://www.sun.com/software/bandwidth))
  - Contains additional product information, such as:
    - Brief overview
    - Features and benefits
    - Software system requirements
    - Product datasheet
    - FAQs
    - White papers
    - Virtual demo



Sun Microsystems, Inc.  
901 San Antonio Road  
Palo Alto, CA 94303

1 (800) 786.7638  
+1.512.434.1511

<http://www.sun.com/solaris>