

WAN Compression FAQs

Document ID: 9289

Questions

Introduction
Compression Overview
Implement Compression in Cisco Routers
Troubleshoot Compression
Related Information

Introduction

This document answers the frequently asked questions (FAQs) about WAN compression. This document includes the Compression Overview, Implement Compression in Cisco Routers, and Troubleshoot Compression sections.

Compression Overview

Q. How does data compression work?

A. Data compression works by the identification of patterns in a stream of data. Data compression chooses a more efficient method to represent the same information. Essentially, an algorithm is applied to the data in order to remove as much redundancy as possible. The efficiency and effectiveness of a compression scheme is measured by its compression ratio, the ratio of the size of uncompressed data to compressed data. A compression ratio of 2:1 (which is relatively common) means that the compressed data is half the size of the original data.

There are many different algorithms available in order to compress data. Some algorithms are designed to take advantage of a specific medium and the redundancies found in them. However, they do a poor job when applied to other sources of data. For example, the Motion Picture Experts Group (MPEG) standard is designed to take advantage of the relatively small difference between one frame and another in video data. It does an excellent job in compression of motion pictures, but does not compress the text well.

One of the most important ideas in compression theory is that there exists a theoretical limit, known as Shannon's Limit. This limit tells you how far you can compress a given source of data. Beyond this point, it is impossible to reliably recover compressed data. Modern compression algorithms coupled with the fast processors available today allow users to approach Shannon's Limit. However, they can never cross it.

Refer to these documents for more information on on Shannon's Limit:

- ◆ Shannon's Law
- ◆ Non-Parametric Approach and Shannon's Superresolution Limit

Q. What is the difference between software and hardware compression? Does hardware compression provide a better compression ratio than software compression?

A. Hardware compression and software compression refer to the site in the router to which the compression algorithm is applied. In software compression, it is implemented in the main CPU as a software process. In hardware compression, the compression computations are offloaded to a secondary hardware module. This frees the central CPU from the computationally intensive task of compression calculations.

If you assume that the router has the available clock cycles in order to perform the compression calculations for example, CPU utilization remains at less than 100 percent then there is no difference in the efficiency of hardware compression or software compression.

The achieved compression ratio is a function of the compression algorithm selected and the amount of redundancy in the data to be compressed. It is not where the compression calculations take place.

Q. What is Layer 2 payload compression?

A. Layer 2 payload compression involves the compression of the payload of a Layer 2 WAN protocol, such as PPP, Frame Relay, High-Level Data Link Control (HDLC), X.25, and Link Access Procedure, Balanced (LAPB). The Layer 2 header is untouched by the act of compression. However, the entire contents of the payload (that include higher-layer protocol headers) are compressed. They are compressed as described in How does data compression work?, and use either a form of the "stacker" algorithm (based on the industry standard Lemple Ziv algorithm; refer to the American National Standards Institute (ANSI) document X3.241-1994), or the "predictor" algorithm, which is an older algorithm that is mostly used in legacy configurations.

Q. What is TCP/IP header compression?

A. TCP/IP header compression removes some of the redundant fields in the header of a TCP/IP connection. Header compression keeps a copy of the original header on either side of the link, removes the entirely redundant fields, and differentially codes the remaining fields in order to allow the compression of 40 bytes of header down to an average of 5 bytes. This uses a very specific algorithm designed around the constant structure of the TCP/IP header. It does not touch the payload of the TCP packet in any way. Refer to RFC 1144, Compressing TCP/IP Headers for Low-Speed Serial Links .

Q. When do I need to use TCP/IP header compression instead of Layer 2 payload compression?

A. TCP/IP header compression is designed to be used for slow serial links of 32 k or less, and to produce a significant performance impact. It requires highly interactive traffic with small packet sizes. In such traffic, the ratio of Layer 3 and Layer 4 header to payload is relatively high. Therefore, performance can be improved if you shrink the headers.

Layer 2 payload compression applies the selected compression algorithm to the entire frame payload, which includes the TCP/IP headers. It is designed to be used on links that operate at speeds from 56 k to 1.544 M. It is useful on all types of traffic, as long as the traffic has not been previously compressed by a higher-layer application.

Q. Can you use both TCP/IP header compression and Layer 2 payload compression at the same time?

A. No. You do *not* implement both Layer 2 payload compression and TCP/IP header compression concurrently because:

- ◆ It is redundant and wasteful.
- ◆ Often, the link does not come up or does not pass IP traffic.

Use only Layer 2 payload compression, rather than both Layer 2 payload compression and TCP/IP header compression.

Implement Compression in Cisco Routers

Q. What software version do I need to run for compression?

A. The most recent release in either the Cisco IOS® Software Release 11.3T or 12.0 (mainline, S, or T) train of code is recommended in order to ensure hardware and software compatibility. In addition, Cisco highly recommends that you run the same version of code on both sides of the WAN link in order to ensure compatibility.

Q. What hardware compression module is right for a particular router?

A. This table shows all routers that support hardware compression and the supported modules:

Router	Hardware Compression Adapter
7200 and 7500	SA-COMP/1= and SA-COMP/4=
3620 and 3640	NM-COMPR=
3660	AIM-COMPR4=
2600	AIM-COMPR2=

Note: The Cisco 7200 VXR series of routers does not support the SA-COMP/1= or the SA-COMP/4=. There is no hardware compression adapter for the 7200 VXR series of routers.

Q. What compression protocols are supported in hardware?

A. Cisco hardware compression adapters only support PPP stacker compression and Frame Relay FRF.9 stacker compression. All compression adapters support both of these protocols. Refer to the Frame Relay Forum Web site, and choose **Implementation Agreements** under the Frame Relay menu for more information on the FRF.9 specification.

Q. When do you need a hardware compression module?

A. There is no simple answer to this question, due to differences in traffic patterns and potential configurations of a given router.

Compression is very processor-intensive, and processor utilization is proportional to the amount of traffic you wish to compress. If the router in question has many

processor-intensive features that already run on it, few clock cycles remain for compression.

Compression also requires memory in order to store the reconstruction dictionaries. Therefore, routers short on memory can run into problems. In a hub-and-spoke configuration, the hub often requires a compression module while the spokes do not.

The only way to answer this question is to suggest that you implement compression in stages, and monitor processor utilization.

Q. What is distributed compression?

A. Distributed compression is available when the interface to be compressed sits in a Versatile Interface Processor 2 (VIP2) slot. The compression calculations then offload into the VIP2 processor.

Q. How do you enable compression in hardware?

A. The router defaults to offloading the compression calculations as far away from the CPU as possible. The entire point of hardware compression is to remove the load from the router CPU and place it on the hardware module. If there is a compression module available, it is used for compression. If a compression module is not available, and if the interface in question resides in a VIP2 slot, then the processor on the VIP2 is used for the compression calculations. If that processor is not available, compression is done in software. Specification of either **software**, **distributed**, or **csa #** at the end of a compression command can force the router to use either the main CPU, the VIP2 CPU, or a hardware module, respectively.

Q. What is the difference between an SA-COMP/1= and an SA-COMP/4=?

A. Both compression service adapters have the same processor onboard. The only difference lies in the onboard memory. They can process the same amount of traffic, in terms of both the amount of data and the packets per second (pps).

A service adapter can process up to 60 Mbps aggregate bi-directional uncompressed bandwidth, with 40,000 pps bi-directional or up to 30,000 pps in one direction. As a rule of thumb, one service adapter can run eight compressed E1s. This assumes a 2:1 compression ratio; a 1.7:1 or 1.8:1 is more common.

A COMP/1 has 768 KB of memory that allows it to support 64 different "contexts".

A COMP/4 has 3 MB of memory that allows it to support 256 different "contexts".

One context is essentially one bi-directional reconstruction dictionary pair, that is, one point-to-point link. So, each Frame Relay point-to-point sub-interface is one context. (More specifically, each individual vc has one context associated with it, as Cisco compression works on a "per data-link connection identifier (DLCI)" basis.)

Q. Does Cisco support compression over multi-link PPP?

A. Multi-link PPP with software compression, which includes multi-link PPP with interleaving plus compression, is supported.

Multi-link PPP with hardware compression is supported from Cisco IOS Software Release 12.0(7)T and 12.0(7) on Cisco 7200 and 3600 routers. However, multi-link PPP and Compression Service Adapter (CSA) are not supported on Cisco 7500 routers.

Q. How do you monitor the compression on a link?

A. Issue the **show compression** command, along with the **show interface** command, in order to determine the throughput, the number of compressed packets, and the compression ratio.

Troubleshoot Compression

Q. What are the common problems in the implementation of compression?

1. Using software Layer 2 Payload compression, Cisco only supports first-in, first-out (FIFO) queuing as the packet is compressed prior to presentation to interface queue. Weighted Fair Queuing is on by default. In order to turn it off you need to issue the **no fair-queue** command.
2. Using hardware Layer 2 Payload compression, Fancy queueing is supported as packets are queued prior to being compressed thus enabling successful classification.
- 3.

Q. When you attempt to run fast switching and compression, the packets are process switched. Why?

A. When you run software compression, all packets must go through the processor anyway, and they are process switched. This is the way compression works.

Q. When you type "show compress", you receive either no response or erroneous responses. Why?

A. Show compress is broken in earlier versions of Cisco IOS Software Release 12.0 code. Upgrade to Cisco IOS Software Release 12.0(7) (mainline, S, or T) for the fix (CSCdk15127 (registered customers only)). This is a cosmetic problem only.

Q. What causes problems when you enable compression between a Cisco and an Ascend router?

A. It is a problem with the default configuration on the Ascend box. Contact your Lucent Technologies Technical Support Representative.

Q. When you run Frame Relay payload-compress FRF.9 stac, some higher-level protocols are compressed, but others are not. Why?

A. This is known problem Cisco bug ID CSCdk39968 (registered customers only) . The solution is to upgrade to Cisco IOS Software Release 11.3(7) or later code.

Q. Why does the show compress command show software compression when hardware compression is enabled?

A. This can happen for a number of reasons:

- ◆ If a link is in a shutdown state, issue the **show compress** command in order to show that it runs software compression. When the link comes up, it shows hardware compression. The command shows this because of the necessity to negotiate hardware compression, either through CCP for PPP, or through the FRF.9 process for Frame Relay. In order to run this negotiation, the link must not be shut down.
- ◆ When you run hardware compression over PPP with some earlier versions of Cisco IOS software, do not type **compress stac** in order to issue the command, it is necessary to type **ppp compress stac** in order to issue the command. This is a holdover from an earlier command syntax.
- ◆ In order to run hardware compression in a 7500 series router, the compression service adapter must be in the same VIP2 as the interface that is to be compressed. Interfaces on other VIP2s and on interface processor cards cannot communicate with the compression service adapters.

Q. What does it mean if show compress indicates that you have a compression ratio of less than one? What can cause it?

A. A compression ratio of less than one means that the compression algorithm increases the size of the data. It does not decrease the size of the data. This is caused by one of these reasons:

- ◆ If you try to compress data that has already gone through a compression algorithm at a higher layer. Compression algorithms are designed with the assumption that there exists redundancy to be removed, and the algorithm performs its calculations accordingly. If data has already been compressed, the redundancy has already been removed, and if you apply another compression algorithm to the same data, it can result in the expansion of the data. Such a result occurs if you try to compress at Layer 2 large data packets that contain zipped data. The only previously uncompressed portion of the payload are the TCP/IP headers. A large data packet (such as FTP) can expand such that the total compression ratio is less than one.
- ◆ Compression ratios of less than one can result from an overly taxed CPU. If you run software compression on a router that does not have the cycles in order to perform the necessary calculations, the process stops. One symptom of this is compression ratios of less than one. The only solutions are to remove compression from some links, or to install a hardware compression module.

Related Information

- Cisco IOS Data Compression
- Data Compression AIM for the Cisco 2600 Series
- Data Compression AIM for the Cisco 3660 Series
- Routing IP, chapter on *Compressing TCP Headers*
- Configuring the SA-COMP/1 and SA-COMP/4 Data Compression Service Adapters
- Technical Support & Documentation – Cisco Systems

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2007 – 2008 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)