

A Negative Acknowledgement Mechanism for Signaling Compression

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document describes a mechanism that allows Signaling Compression (SigComp) implementations to report precise error information upon receipt of a message which cannot be decompressed. This negative feedback can be used by the recipient to make fine-grained adjustments to the compressed message before retransmitting it, allowing for rapid and efficient recovery from error situations.

Table of Contents

1. Introduction	2
1.1. The Problem	2
1.1.1. Compartment Disposal	3
1.1.2. Client Restart	3
1.1.3. Server Failover	3
1.2. The Solution	4
2. Node Behavior	4
2.1. Normal SigComp Message Transmission	4
2.2. Receiving a "Bad" SigComp Message	5
2.3. Receiving a SigComp NACK	6
2.3.1. Unreliable Transport	6
2.3.2. Reliable Transport	6
2.4. Detecting Support for NACK	7
3. Message Format	7
3.1. Message Fields	8
3.2. Reason Codes	9
4. Security Considerations	13
4.1. Reflector Attacks	13
4.2. NACK Spoofing	13
5. IANA Considerations	14
6. Acknowledgements	14
7. References	14
7.1. Normative References	14
7.2. Informative References	14

1. Introduction

Signaling Compression [1], often called "SigComp", defines a protocol for transportation of compressed messages between two network elements. One of the key features of SigComp is the ability of the sending node to request that the receiving node store state objects for later retrieval.

1.1. The Problem

While the "SigComp - Extended Operations" document [2] defines a mechanism that allows for confirmation of state creation, operational experience with the SigComp protocol has demonstrated that there are still several circumstances in which a sender's view of the shared state differs from the receiver's view. A non-exhaustive list detailing the circumstances in which such failures may occur is below.

1.1.1. Compartment Disposal

In SigComp, stored states are associated with compartments. Conceptually, the compartments represent one instance of a remote application. These compartments are used to limit the amount of state that each remote application is allowed to store. Compartments are created upon receipt of a valid SigComp message from a remote application. In the current protocol, applications are expected to signal when they are finished with a compartment so that it can be deleted (by using the S-bit in requested feedback data).

Unfortunately, expecting the applications to be well-behaved is not sufficient to prevent state from piling up. Unexpected client failures, reboots, and loss of connectivity can cause compartments to become "stuck" and never removed. To prevent this situation, it becomes necessary to implement a scheme by which compartments that appear disused may eventually be discarded.

While the preceding facts make such a practice necessary, discarding compartments without explicit signaling can have the unfortunate side effect that active compartments are sometimes discarded. This leads to a different view of state between the server and the client.

1.1.2. Client Restart

The prime motivation for SigComp was compression of messages to be sent over a radio interface. Consequently, most deployments of SigComp will involve a mobile unit as one of the endpoints. Mobile terminals are generally not guaranteed to be available for extended durations of time. Node restarts (due to, for example, a battery running out) will induce situations in which the network-based server believes that the client contains several states that are no longer actually available.

1.1.3. Server Failover

Many applications for which SigComp will be used (e.g., SIP [3]) use DNS SRV records for server lookup. One of the important features of DNS SRV records is the ability to specify multiple servers from which clients will select at random, with probabilities determined by the q-value weighting. The reason for defining this behavior for SRV records is to allow load distribution through a set of equivalent servers, and to permit clients to continue to function even if the server with which they are communicating fails. When using protocols that use SRV for such distribution, the traffic to a failed server is typically sent by the client to an equivalent server that can serve

the same purpose. From an application perspective, this new server often appears to be the same endpoint as the failed server, and will consequently resolve to the same compartment.

Although SigComp state can be replicated amongst such a cluster of servers, maintaining integrity of such states requires a two-phase commit process that adds a great deal of complexity to the server and can degrade performance significantly.

1.2. The Solution

Although SigComp allows returned SigComp parameters to signal that all states have been lost (by setting "state_memory_size" to 0 for one message in the reverse direction), such an approach provides an incomplete solution to the problem. In addition to wiping out an entire compartment when only one state is corrupt or missing, this approach suffers from the unfortunate behavior that it requires a message in the reverse direction that the remote application will authorize. Unless a lower-layer security mechanism is employed (e.g., TLS), this would typically mean that a compressed application-level message in the reverse direction must be sent before recovery can occur. In many cases (such as SIP-based mobile terminals), these messages won't be sent often; in others (pure client/server deployments), they won't ever be sent.

The proposed solution to this problem is a simple Negative Acknowledgement (NACK) mechanism which allows the recipient to communicate to the sender that a failure has occurred. This NACK contains a reason code that communicates the nature of the failure. For certain types of failures, the NACK will also contain additional details that might be useful in recovering from the failure.

2. Node Behavior

The following sections detail the behavior of nodes sending and receiving SigComp NACKs. The actual format and values are described in Section 3.

2.1. Normal SigComp Message Transmission

Although normal in all other respects, SigComp implementations that use the NACK mechanism need to calculate and store a SHA-1 hash for each SigComp message that they send. This must be stored in such a way that, given the SHA-1 hash, the implementation is able to locate the compartment with which the sent message was associated.

In other words, if someone hands the SHA-1 hash back to the compressor, it needs to be able to find the compartment with which it was working when it sent the message with that hash. This only requires that the compressor knows with which compartment it is working when it sends a message (which is always the case), and that the SHA-1 hash, when stored, points to that compartment in some way.

2.2. Receiving a "Bad" SigComp Message

When a received SigComp message causes a decompression failure, the recipient forms and sends a SigComp NACK message. This NACK message contains a SHA-1 hash of the received SigComp message that could not be decompressed. It also contains the exact reason decompression failed, as well as any additional details that might assist the NACK recipient to correct any problems. See Section 3 for more information about formatting the NACK message and its fields.

For a connection-oriented transport, such as TCP, the NACK message is sent back to the originator of the failed message over that same connection.

For a stream-based transport, such as TCP, the standard SigComp delimiter of 0xFFFF is used to terminate the NACK message.

For a connectionless transport, such as UDP, the NACK message is sent back to the originator of the failed message at the port and IP address from which the message was sent. Note that this may or may not be the same port on which the application would typically receive messages. To accommodate implementations that use connect() or similar constructs, the NACK will be sent from the IP address and port to which the uninterpretable message was sent. From a practical perspective, this is probably easiest to determine by binding listening sockets to a specific interface; however, other mechanisms may also be employed.

The behavior specified above is strictly necessary for any generally useful form of a NACK mechanism. In the most general case, when an implementation receives a message that it cannot decompress, it has exactly three useful pieces of information: (1) the contents of the message, (2) an indication of why the message cannot be decoded, and (3) the IP address and port from which the message originated. Note that none of these contains any indication of where the remote application is listening for messages, if it differs from the sending port.

2.3. Receiving a SigComp NACK

The first action taken upon receipt of a NACK is an attempt to find the message to which the NACK corresponds. This search is performed using the 20-byte SHA-1 hash contained in the NACK. Once the matching message is located, further operations are performed based on the compartment that was associated with the sent message.

Further behavior of a node upon receiving a SigComp NACK depends on whether a reliable or unreliable transport is being used.

2.3.1. Unreliable Transport

When SigComp is used over an unreliable transport, the application has no reasonable expectation that the transport layer will deliver any particular message. It then becomes the application layer's responsibility to ensure that data is retransmitted as necessary. In these circumstances, the NACK mechanism relies on such behavior to ensure delivery of the message, and never performs retransmissions on the application's behalf.

When a NACK is received for a message sent over an unreliable transport, the NACK recipient uses the contained information to make appropriate adjustments to the compressor associated with the proper compartment. The exact nature of these adjustments are specific to the compression scheme being used, and will vary from implementation to implementation. The only requirement on these adjustments is that they must have the effect of compensating for the error that has been indicated (e.g., by removing the state that the remote node indicates it cannot retrieve).

In particular, when an unreliable transport is used, the original message must not be retransmitted by the SigComp layer upon receipt of a NACK. Instead, the next application-initiated transmission of a message will take advantage of the adjustments made as a result of processing the NACK.

2.3.2. Reliable Transport

When a reliable transport is employed, the application makes a basic assumption that any message passed down the stack will be retransmitted as necessary to ensure that the remote node receives it, unless a failure is indicated by the transport layer. Because SigComp acts as a shim between the transport-layer and the application, it becomes the responsibility of the SigComp implementation to ensure that any failure to transmit a message is communicated to the application.

When a NACK is received for a message sent over a reliable transport, the SigComp layer must indicate to the application that an error has occurred. In general, the application should react in the same way as it does for any other transport layer error, such as a TCP connection reset. For most applications, this reaction will initially be an attempt to reset and re-establish the connection, and re-initiate the failed transaction. The SigComp layer should also use the information contained in the NACK to make appropriate adjustments to the compressor associated with the proper compartment (similar to the adjustments made for unreliable transport). Thus, if the compartment is not reset by resetting the TCP connection, the next message will take advantage of the adjustments.

2.4. Detecting Support for NACK

Detection of support for the NACK mechanism may be beneficial in certain circumstances. For example, with the current definition of SigComp, acknowledgment of state receipt is required before a sender can reference such state. When multiple messages are sent before a response is received, the need to wait for such responses can cause significant decreases in message compression efficiency. If it is known that the receiver supports the NACK mechanism, the sender can instead optimistically assume that the state created by a sent message has been created, and is allowed to be referenced. If such an assumption turns out to be false (due to, for example, packet loss or packet reordering), the sender can recover upon receipt of a NACK.

In order to facilitate such detection, any implementation that will send NACK messages upon decompression failure will indicate a SigComp version number of 0x02 in its Universal Decompressor Virtual Machine (UDVM). The bytecodes sent to such an endpoint can check the version number, and send appropriate indication back to their compressor as requested feedback. Except for the NACK mechanism described in this document, implementations advertising a version of 0x02 behave exactly like those advertising a version number of 0x01.

3. Message Format

SigComp NACK packets are syntactically valid SigComp messages which have been specifically designed to be safely ignored by implementations that do not support the NACK mechanism.

In particular, NACK messages are formatted as the second variant of a SigComp message (typically used for code upload) with a "code_len" field of zero. The NACK information (message identifier, reason for failure, and error details) is encoded in the "remaining SigComp

message" area, typically used for input data. Further, the "destination" field is used as a version identifier to indicate which version of NACK is being employed.

3.1. Message Fields

The format of the NACK message and the use of the fields within it are shown in Figure 1.

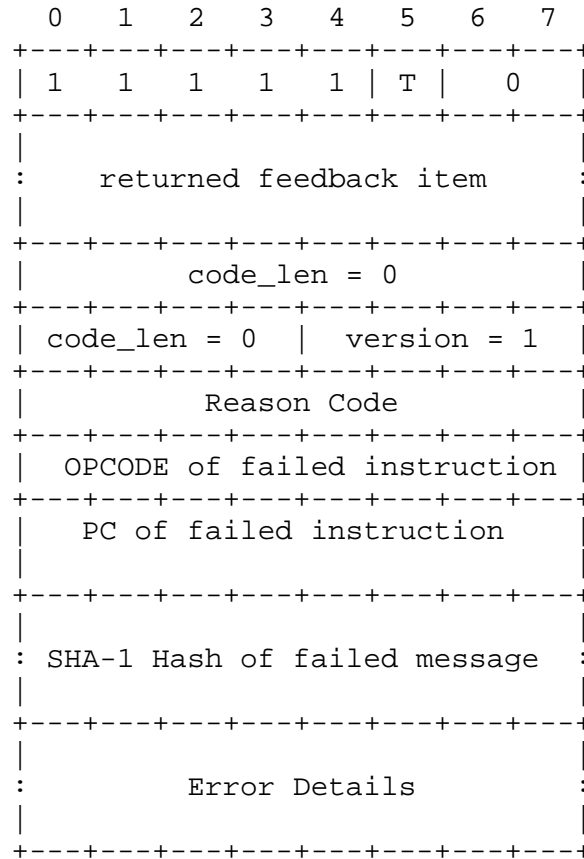


Figure 1: SigComp NACK Message Format

- o "Reason Code" is a one-byte value that indicates the nature of the decompression failure. The specific codes are given in Section 3.2.
- o "OPCODE of failed instruction" is a one-byte field that includes the opcode to which the PC was pointing when the failure occurred. If failure occurred before the UDVM began executing any code, this field is set to 0.

- o "PC of failed instruction" is a two-byte field containing the value of the program counter when failure occurred (i.e., the memory address of the failed UDVM instruction). The field is encoded with the most significant byte of the PC first (i.e., in network or big endian order). If failure occurred before the UDVM began executing any code, this field is set to 0.
- o "SHA-1 Hash of failed message" contains the full 20-byte SHA-1 hash of the SigComp message that could not be decompressed. This information allows the NACK recipient to locate the message that failed to decompress so that adjustments to the correct compartment can be performed. When performing this hash, the entire SigComp message is used, from the header byte (binary 11111xxx) to the end of the input. Any lower-level protocol headers (such as UDP or IP) and message delimiters (the 0xFFFF that marks message boundaries in stream protocols) are not included in the hash. When used over a stream based protocol, any 0xFFxx escape sequences are un-escaped before performing the hash operation.
- o "Error Details" provides additional information that might be useful in correcting the problem that caused decompression failure. Its meaning is specific to the "Reason Code". See Section 3.2 for specific information on what appears in this field.
- o "Code_len" is the "code_len" field from a standard SigComp message. It is always set to "0" for NACK messages.
- o "Version" gives the version of the NACK mechanism being employed. This document defines version 1.

3.2. Reason Codes

Note that many of the status codes are more useful in debugging interoperability problems than with on-the-fly correction of errors. The "STATE_NOT_FOUND" error is a notable exception: it will generally cause the NACK recipient to encode future messages so as to not use the indicated state.

Upon receiving the other status messages, an implementation would typically be expected either to use a different set of bytcodes or, if that is not an option, to send that specific message uncompressed.

Error	Code	Details
STATE_NOT_FOUND	1	State ID (6 - 20 bytes)
CYCLES_EXHAUSTED	2	Cycles Per Bit (1 byte)
USER_REQUESTED	3	
SEGFault	4	
TOO_MANY_STATE_REQUESTS	5	
INVALID_STATE_ID_LENGTH	6	
INVALID_STATE_PRIORITY	7	
OUTPUT_OVERFLOW	8	
STACK_UNDERFLOW	9	
BAD_INPUT_BITORDER	10	
DIV_BY_ZERO	11	
SWITCH_VALUE_TOO_HIGH	12	
TOO_MANY_BITS_REQUESTED	13	
INVALID_OPERAND	14	
HUFFMAN_NO_MATCH	15	
MESSAGE_TOO_SHORT	16	
INVALID_CODE_LOCATION	17	
BYTECODES_TOO_LARGE	18	Memory size (2 bytes)
INVALID_OPCODE	19	
INVALID_STATE_PROBE	20	
ID_NOT_UNIQUE	21	State ID (6 - 20 bytes)
MULTILOAD_OVERWRITTEN	22	
STATE_TOO_SHORT	23	State ID (6 - 20 bytes)
INTERNAL_ERROR	24	
FRAMING_ERROR	25	

Only the five errors "STATE_NOT_FOUND", "CYCLES_EXHAUSTED", "BYTECODES_TOO_LARGE", "ID_NOT_UNIQUE", and "STATE_TOO_SHORT" contain details; for all other error codes, the "Error Details" field has zero length.

Figure 2: SigComp NACK Reason Codes

1. STATE_NOT_FOUND

A state that was referenced cannot be found. The state may have been referenced by the UDVM executing a STATE-ACCESS instruction; it also may have been referenced by the "partial state identifier" field in a SigComp message. The "details" field contains the state identifier for the state that could not be found. This is also the proper error to return in the case that a unique state item was matched but fewer bytes of state ID were sent than required by the `minimum_access_length`.

2. CYCLES_EXHAUSTED
Decompression of the message has taken more cycles than were allocated to it. The "details" field contains a one-byte value that communicates the number of cycles per bit. The cycles per bit is represented as an unsigned 8-bit integer (i.e., not encoded).
3. USER_REQUESTED
The DECOMPRESSION-FAILURE opcode has been executed.
4. SEGFAULT
An attempt to read from or write to memory that is outside of the UDVM's memory space has been attempted.
5. TOO_MANY_STATE_REQUESTS
More than four requests to store or delete state objects have been requested.
6. INVALID_STATE_ID_LENGTH
A state id length less than 6 or greater than 20 has been specified.
7. INVALID_STATE_PRIORITY
A state priority of 65535 has been specified when attempting to store a state.
8. OUTPUT_OVERFLOW
The decompressed message is too large to be decoded by the receiving node.
9. STACK_UNDERFLOW
An attempt to pop a value off the UDVM stack was made with a stack_fill value of 0.
10. BAD_INPUT_BITORDER
An INPUT-BITS or INPUT-HUFFMAN instruction was encountered with the "input_bit_order" register set to an invalid value (i.e., one of the upper 13 bits is set).
11. DIV_BY_ZERO
A DIVIDE or REMAINDER opcode was encountered with a divisor of 0.
12. SWITCH_VALUE_TOO_HIGH
The input to a SWITCH opcode exceeds the number of branches defined.

13. `TOO_MANY_BITS_REQUESTED`
An `INPUT-BITS` or `INPUT-HUFFMAN` instruction was encountered that attempted to input more than 16 bits.
14. `INVALID_OPERAND`
An operand for an instruction could not be resolved to an integer value (e.g., a literal or reference operand beginning with 11111111).
15. `HUFFMAN_NO_MATCH`
The input string does not match any of the bitcodes in the `INPUT-HUFFMAN` opcode.
16. `MESSAGE_TOO_SHORT`
When attempting to decode a SigComp message, the recipient determined that there were not enough bytes in the message for it to be valid.
17. `INVALID_CODE_LOCATION`
The "code location" field in the SigComp message was set to the invalid value of 0.
18. `BYTECODES_TOO_LARGE`
The bytecodes that a SigComp message attempted to upload exceed the amount of memory available in the receiving UDVM. The details field is a two-byte expression of the `DECOMPRESSION_MEMORY_SIZE` of the receiving UDVM. This value is communicated most-significant-byte first.
19. `INVALID_OPCODE`
The UDVM attempted to identify an undefined byte value as an instruction.
20. `INVALID_STATE_PROBE`
When attempting to retrieve state, the `state_length` operand is set to 0 but the `state_begin` operand is non-zero.
21. `ID_NOT_UNIQUE`
A partial state identifier that was used to access state matched more than one state item. Note that this error might be returned as the result of executing a `STATE-ACCESS` instruction or attempting to locate a unique piece of state as identified by the "partial state identifier" in a SigComp message. The "details" field contains the partial state identifier that was requested.
22. `MULTILOAD_OVERWRITTEN`
A `MULTILOAD` instruction attempted to overwrite itself.

23. STATE_TOO_SHORT

A STATE-ACCESS instruction has attempted to copy more bytes from a state item than the state item actually contains. The "details" field contains the partial state identifier that was requested. Implementors are cautioned to return only the partial state identifier that was requested; if the NACK contains any state identifier in addition to what was requested, attackers may be able to use that additional information to access the state.

24. INTERNAL_ERROR

The UDVM encountered an unexpected condition that prevented it from decompressing the message.

25. FRAMING_ERROR

The UDVM encountered a framing error (unquoted 0xFF 80 .. 0xFF FE in an input stream.) This error is applicable only to messages received on a stream transport. In the case of a framing error, a SHA-1 hash for a unique message cannot be determined. Consequently, when a FRAMING_ERROR NACK is sent, the "SHA-1 Hash of failed message" field should be set to all zeros.

4. Security Considerations

4.1. Reflector Attacks

Because SigComp NACK messages are by necessity sent in response to other messages, it is possible to trigger them by intentionally sending malformed messages to a SigComp implementation with a spoofed IP address. However, because such actions can only generate one message for each message sent, they don't serve as amplifier attacks. Further, due to the reasonably small size of NACK packets, there cannot be a significant increase in the size of the packet generated.

It is worth noting that nearly all deployed protocols exhibit this same behavior.

4.2. NACK Spoofing

Although it is possible to forge NACK messages as if they were generated by a different node, the damage that can be caused is minimal. Reporting a loss of state will typically result in nothing more than the re-transmission of that state in a subsequent message. Other failure codes would result in the next message being sent using an alternate compression mechanism, or possibly uncompressed.

Although all of the above consequences result in slightly larger messages, none of them have particularly catastrophic implications for security.

5. IANA Considerations

This document defines a new value for the IANA registered attribute SigComp_version.

Value (in hex): 02

Description: SigComp version 2 (NACK support)

Reference: [RFC4077]

6. Acknowledgements

Thanks to Carsten Bormann, Zhigang Liu, Pekka Pessi, and Robert Sugar for their comments and suggestions. Special thanks to Abigail Surtees and Richard Price for several very detailed reviews and suggestions.

7. References

7.1. Normative References

- [1] Price, R., Bormann, C., Christoffersson, J., Hannu, H., Liu, Z., and J. Rosenberg, "Signaling Compression (SigComp)", RFC 3320, January 2003.
- [2] Hannu, H., Christoffersson, J., Forsgren, S., Leung, K.-C., Liu, Z., and R. Price, "Signaling Compression (SigComp) - Extended Operations", RFC 3321, January 2003.

7.2. Informative References

- [3] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

Author's Address

Adam Roach
Estacado Systems
17210 Campbell Road
Suite 250
Dallas, TX 75252
US

EMail: adam@estacado.net

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

