Q. What is IF-M Segmentation?

A. IF-M Segmentation provides a standard means to manage the size of IF-M messages between TNC clients and servers. It also provides a mechanism by which large messages can be delivered in segments, to avoid overwhelming the network connection or the memory capacity of either the client or the server. This feature can be especially useful for tools which use the TNC architecture to deliver large volumes of data.

Q. How does IF-M Segmentation control maximum message size?

A. IF-M Segmentation supports the creation of a "Segmentation Contract". Both the data sender and data recipient must explicitly agree to the Segmentation Contract to be bound by it. Once in place, the Segmentation Contract imposes an upper limit on the size of messages that the contracted data sender can send to the contracted data recipient. If a message exceeds that size limit, the data sender instead sends an error message to alert the data sender to the presence of the overlarge message, enabling the data recipient to take the appropriate response. This can include granting a one-time exemption from the Segmentation Contract to allow the overlarge message to be transmitted. Alternatively, the data recipient could chose not to have the overlarge message delivered.

Q. How does IF-M Segmentation facilitate segmented delivery of large attributes?

A. The Segmentation Contract identifies the largest message that can be sent whole. If a message is larger than that size, but still smaller than the maximum permitted message size, then the data sender employs a Segmented Message Exchange. In a Segmented Message Exchange, the overlarge message is broken into segments, which are transmitted one at a time. The data recipient must explicitly request each segment. This allows the data recipient to control the rate of delivery of the segments. It also gives the data recipient the option of terminating the Segmented Message Exchange at any time. For example, this might happen if the specific question the data recipient is hoping to answer is answered early in the delivered message, or if changes to the situation on the data recipient obviate the need for the information.

Q. Isn't this capability already provided by IP packet fragmentation?

A. IP packet fragmentation happens entirely at the network layer, breaking up large IP datagrams to allow their transmission over connections with smaller "maximum transmission units" (MTUs). . In most cases, the datagram is not made available to the receiving software until the entire message has been delivered, and the message still ends up requiring a large chunk of memory after reassembly. By contrast, IF-M Segmentation is controlled at the application layer. Specifically, the TNC architecture components responsible for sending and receiving IF-M interface messages explicitly control the delivery of segments. This allows these components to manage transmission rates to fit with their own processing abilities, or even terminate exchanges that are no longer needed.

Q. What does IF-M Segmentation offer to users of Trusted Network Communications (TNC)?

A. IF-M Segmentation allows components communicating over IF-M to directly manage the size and rate of data that is exchanged between them. The TNC architecture allows IF-M messages of up to 4 GB in size, but messages that large can consume significant network and memory resources. IF-M Segmentation allows smaller message limits to be imposed, and allows more managed delivery of larger messages to avoid flooding the network or overwhelming the memory of a message sender or recipient.

An analogy would be choosing to stream content, rather than download it; downloading requires retrieving the entire content at once and waiting until it's complete, whereas streaming enables retrieving the content serially and using portions of it as it arrives.

Q. What IF-M exchanges can use IF-M Segmentation?
A. All of them; IF-M Segmentation is agnostic as to the messages it conveys, and any component supporting IF-M messages can add support for IF-M Segmentation.

Q. Can Server Discovery and Validation be used with IETF's Network Endpoint Assessment (NEA)?
A. Yes. NEA and TNC are designed to be interoperable with each other. IF-M Segmentation can be used to manage exchanges between NEA Posture Collectors and Posture Validators.

Q. Who has implemented IF-M Segmentation?
A. IF-M Segmentation has been integrated into strongSwan 5.2.1 and later.