TCP SACK Panic
June 2019

CVSSv3 Score 7.5

Important

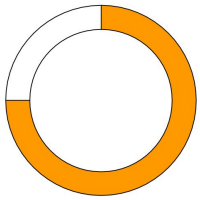

Advisory ID	nutanix-sa-015-sack	CVE(s)
		CVE-2019-11477 CVE-2019-11478 CVE-2019-11479
Last Updated	23 October 2019	
Published	18 June 2019	
Version	10.0	

Final Update

Summary

Originating from a Netflix security advisory (see sources below) they discovered several TCP networking vulnerabilities in FreeBSD and Linux kernels. These vulnerabilities relate to maximum segment size (MSS) and TCP Selective Acknowledgement (SACK). The more critical of the three vulnerabilities, also known as SACK Panic, can allow for a remotely triggered kernel panic in Linux kernels.

Information on Vulnerabilities

Three vulnerabilities, outlined below, make up this collective advisory with CVE-2019-11477 (SACK Panic) being the most critical. Definition of two terms will help clarify the issue.

MSS - Maximum Segment Size

The maximum segment size is a parameter set in the TCP header of a packet that specifies the total data contained within a TCP segment. This information is necessary in situations where packets become fragmented as they are transmitted across different routes. This parameter informs the receiving host how large the TCP segment size, which is necessary in order to adequately reassemble the packet in the event of fragmentation.

TCP SACKs

During TCP communication, Sequence Numbers (SEQ) and Acknowledgement Numbers (ACK) are used by the client and server to determine which segments have been sent to the client, and which segments the client acknowledges were received. The absence of an ACK for a particular segment during communication would trigger the server to retransmit all segments after the last received segment number.

In order to make this process more efficient Selective Acknowledgement (SACK) was devised as part of RFC-2018. The SACK mechanism allows for the client, in the above example, to explicitly

state exactly which segments it's missing. In this scenario the server only needs to retransmit the missing segments, and not every segment after the last ACK.

CVE-2019-11477 - SACK Panic

This vulnerability relies on a flaw within the Linux kernel where the MSS of a connection is set to its lowest limit of 48 bytes, which only leaves 8 bytes of data per segment. In this scenario, a specially crafted SACK can trigger a denial of service by way of a kernel panic by way of overflowing the tcp_gso_segs parameter in the kernel's Socket Buffers (SKB).

CVE-2019-11478 - SACK Slowness

This vulnerability relies on a resource consumption flaw in the Linux kernel Socket Buffer (SKB) around TCP Selective Acknowledgment (SACK) segments. Specially crafted SACK segments can be sent causing the SKB to become fragmented. This fragmentation leads to increased resource utilization due to the processing of these fragments. As additional SACK segments come in, further fragmentation occurs, eventually resulting in a Denial of Service.

CVE-2019-11479 - Excessive Resource Consumption due to Low MSS Values

This vulnerability relies on setting the MSS of a TCP connection to its lowest value, 48 bytes, which leaves only 8 bytes for actual data on the segment. This low amount of data in the segment results in increased CPU and Memory utilization on the host due to the larger number of segments that must be created to complete the transfer of data. This can result in a Denial of Service by repeatedly sending the server requests with the minimum MSS size of 48 bytes.

Affected Products

This document will be updated with the patch release schedules. Please check the Nutanix Support Portal - Security Advisories for the latest update.

Nutanix Products

Product	Fix Release
AHV	*Now available via 20170830.299 within 5.11 and 5.10.6 on the Nutanix Support Portal.*
Nutanix AOS	*Now available via AOS 5.10.7 and 5.11.1 on the Nutanix Support Portal.*
Prism Central	*Now available via PC 5.11.1 on the Nutanix Support Portal.*
Files	*Now available via Files 3.6 on the Nutanix Support Portal*
Move	*Now available via Move 3.2.0 on the Nutanix Support Portal.*
X-Ray	*Now available via X-Ray 3.6 on the Nutanix Support Portal.*

Era	*Now available via Era 1.1.1.1 on the Nutanix Support Portal.*

Mitigations

Mitigations are possible to these attacks, but at this time we do not recommend implementing these mitigations until proper validation has taken place. If this validation is successful, and updated advisory will be posted with instructions. Final resolution of the above issues is accomplished by a kernel patch.

Sources

RedHat TCP SACK Panic - https://access.redhat.com/security/vulnerabilities/tcpsack
Netflix Original Disclosure - https://github.com/Netflix/security-bulletins/blob/master/advisories/third-party/2019-001.md

Support

If you have questions, please open a case with Nutanix Support at http://portal.nutanix.com, or by calling Support at the phone number on the website http://www.nutanix.com/support.

Thank you for being a Nutanix customer.

Revision History

Version	Section	Date
1.0	-	18 June 2019
2.0	Updated releases and timelines	25 June 2019
3.0	Updated releases and timelines	10 July 2019
4.0	Updated timelines	25 July 2019
5.0	Updated releases	16 August 2019
6.0	Updated Releases	04 September 2019
7.0	Updated timelines	18 September 2019
8.0	Updated timelines	03 October 2019
9.0	Updated releases	08 October 2019
10.0	Final Release	23 October 2019