

mDNSResponder TCP Keepalive (TCPKA)

As mentioned in mDNSResponder Offload.pdf, TCP keepalive (TCPKA) configuration is passed down to the NIC FW via mDNS resource records. However, that document lacks specific details needed for the implementer to understand how to parse the associated resource record.

The format of the TCPKA configuration is dependent on if an IPv4 or an IPv6 keepalive is being passed down. There is no way for the firmware to know if it is dealing with IPv4 or IPv6 until it parses the content of the configuration:

```
RDATA[0]    : length of keepalive text which follows
RDATA[1...] : keepalive text = "t=%d i=%d c=%d h=<IPv4> d=%#a l=%u r=%u m=<MAC> s=%u a=%u w=%u"
RDATA[1...] : keepalive text = "t=%d i=%d c=%d H=<IPv6> D=%#a l=%u r=%u m=<MAC> s=%u a=%u w=%u"
```

Note that IPv4 and IPv6 are different based on 'h=d=' vs. 'H=D='

The “**Timeout**” field (t=%d) shall have units of seconds. This field represents the expected time between successful KA.

The “**Interval**” field (i=%d) shall have units of milliseconds. This field represents the time at which retries should be attempted between un-acknowledged transmissions. In addition, the NIC shall apply a geometric delay between subsequent retries.

The “**Count**” field (c=%d) has no associated units. This field represents the number of unsuccessful KA before the connection is considered broken. After the connection is considered broken, the NIC shall wake the system Host.

For an IPv4 field (h=<IPv4>), the IPv4 address shall be decoded from the format string as : %d.%d.%d.%d.

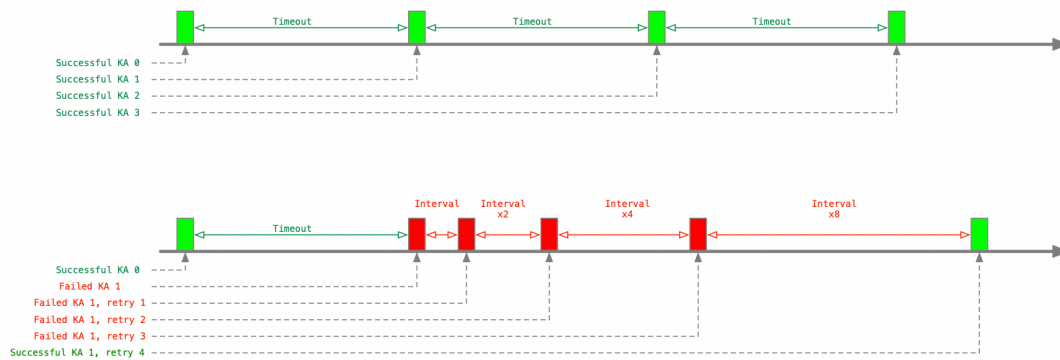
For an IPv6 field (H=<IPv6>), the IPv6 address shall be decoded according to RFC5952 (<https://datatracker.ietf.org/doc/html/rfc5952>).

For the MAC address field (m=<MAC>), the MAC address shall be decoded from the format string as %02X:%02X:%02X:%02X:%02X:%02X.

The following diagrams aim to illustrate how the “**Timeout**” and “**Interval**” fields are related.

The first diagram shows a ‘normal’ case where the NIC’s TCPKA partner is responding normally. After receiving a TCPKA response, the NIC waits for “**Timeout**” before it attempts the next transmission.

The second diagram shows a case where the NIC was unable to get a TCPKA response temporarily. Before each of the subsequent retries, it applies a geometrically scaled “**Interval**”. Finally, a successful response is received on the 5th attempt, and we are back to in the ‘normal’ condition. In this example, if “**Count**” were 4, the NIC would have been expected to wake the system, rather than attempt another retry at the expiration of **Interval x8**.



Sample mDNS RR buffer with IPv4 Keepalive record:

```
astalos NCM-working-group(:|✓) % xxd -g 1 tcpka.mdns
```

```
00000000: 00 00 00 00 00 00 00 00 00 00 00 05 64 75 74 .....dut
00000010: 2d 78 04 5f 73 73 68 04 5f 74 63 70 05 6c 6f 63 -x._ssh._tcp.loc
00000020: 61 6c 00 00 10 80 01 00 00 11 94 00 01 00 09 5f al....._
00000030: 73 65 72 76 69 63 65 73 07 5f 64 6e 73 2d 73 64 services._dns-sd
00000040: 04 5f 75 64 70 c0 1c 00 0c 00 01 00 00 11 94 00 ._udp.....
00000050: 02 c0 12 c0 12 00 0c 00 01 00 00 11 94 00 02 c0 .....
00000060: 0c 05 64 75 74 2d 78 0c 5f 64 65 76 69 63 65 2d ..dut-x._device-
00000070: 69 6e 66 6f c0 17 00 10 00 01 00 00 11 94 00 23 info.....#
00000080: 0e 6d 6f 64 65 6c 3d 4d 61 63 31 34 2c 31 32 0a .model=Mac14,12.
00000090: 6f 73 78 76 65 72 73 3d 32 35 08 69 63 6f 6c 6f osxvers=25.icolo
000000a0: 72 3d 30 05 64 75 74 2d 78 09 5f 73 66 74 70 2d r=0.dut-x._sftp-
000000b0: 73 73 68 c0 17 00 10 00 01 00 00 11 94 00 01 00 ssh.....
000000c0: c0 2e 00 0c 00 01 00 00 11 94 00 02 c0 a9 c0 a9 .....
000000d0: 00 0c 00 01 00 00 11 94 00 02 c0 a3 05 64 75 74 .....dut
000000e0: 2d 78 04 5f 72 66 62 c0 17 00 10 80 01 00 00 11 -x._rfb.....
000000f0: 94 00 01 00 c0 2e 00 0c 00 01 00 00 11 94 00 02 .....
00000100: c0 e2 c0 e2 00 0c 00 01 00 00 11 94 00 02 c0 dc .....
00000110: c0 dc 00 21 80 01 00 00 11 94 00 0e 00 00 00 00 ...!......
00000120: 17 0c 05 64 75 74 2d 78 c0 1c 09 5f 6b 65 72 62 ...dut-x..._kerb
00000130: 65 72 0f 73 c1 22 00 10 00 01 00 00 11 94 00 33 eros.".....3
00000140: 32 4c 4b 44 43 3a 53 48 41 31 2e 42 42 30 39 31 2LKDC:SHA1.BB091
00000150: 41 41 43 35 37 32 38 43 37 41 46 46 39 31 34 46 AAC5728C7AFF914F
00000160: 34 45 32 35 34 36 33 41 31 37 39 44 34 32 32 39 4E25463A179D4229
00000170: 32 43 46 c0 0c 00 21 80 01 00 00 11 94 00 08 00 2CF...!......
00000180: 00 00 00 00 16 c1 22 c0 a3 00 21 80 01 00 00 11 .....".!.....
00000190: 94 00 08 00 00 00 00 00 16 c1 22 05 31 35 35 31 .....".1551
000001a0: 34 0a 5f 6b 65 65 70 61 6c 69 76 65 c0 38 00 0a 4._keepalive.8..
000001b0: 80 01 00 00 11 94 00 73 72 74 3d 35 20 69 3d 33 .....srt=5 i=3
000001c0: 30 20 63 3d 31 30 20 68 3d 31 39 32 2e 31 36 38 0 c=10 h=192.168
000001d0: 2e 38 31 2e 32 35 20 64 3d 31 39 32 2e 31 36 38 .81.25 d=192.168
000001e0: 2e 38 31 2e 33 33 20 6c 3d 35 31 32 35 38 20 72 .81.33 l=51258 r
000001f0: 3d 34 38 30 30 34 20 6d 3d 30 30 3a 33 30 3a 39 =48004 m=00:30:9
00000200: 33 3a 31 30 3a 33 33 3a 34 38 20 73 3d 32 34 34 3:10:33:48 s=244
00000210: 30 37 31 35 33 39 20 61 3d 34 31 38 37 33 32 32 071539 a=4187322
00000220: 31 32 34 20 77 3d 32 30 35 39 00 124 w=2059.
```

Example python script output to parse the resource record:

```
astalos NCM-working-group(:|✓) % python3 parse-rr-bin.py ./tcpka.mdns 15
idx name type class ttl rd_len
0 dut-x._ssh._tcp.local. TXT 0x8001 4500 0x0001
1 _services._dns-sd._udp.local. PTR 0x0001 4500 0x0002
2 _ssh._tcp.local. PTR 0x0001 4500 0x0002
3 dut-x._device-info._tcp.local. TXT 0x0001 4500 0x0023
4 dut-x._sftp-ssh._tcp.local. TXT 0x8001 4500 0x0001
5 _services._dns-sd._udp.local. PTR 0x0001 4500 0x0002
6 _sftp-ssh._tcp.local. PTR 0x0001 4500 0x0002
7 dut-x._rfb._tcp.local. TXT 0x8001 4500 0x0001
8 _services._dns-sd._udp.local. PTR 0x0001 4500 0x0002
9 _rfb._tcp.local. PTR 0x0001 4500 0x0002
10 dut-x._rfb._tcp.local. SRV 0x8001 4500 0x000e SRV prio 0, weight 0, port 5900, target dut-x.local.
11 _kerberos.dut-x.local. TXT 0x0001 4500 0x0033
12 dut-x._ssh._tcp.local. SRV 0x8001 4500 0x0008 SRV prio 0, weight 0, port 22, target dut-x.local.
13 dut-x._sftp-ssh._tcp.local. SRV 0x8001 4500 0x0008 SRV prio 0, weight 0, port 22, target dut-x.local.
14 15514._keepalive._dns-sd._udp.local. NULL 0x8001 4500 0x0073 t=5 i=30 c=10 h=192.168.81.25 d=192.168.81.33 l=51258 r=48004 m=00:30:93:10:33:48 s=244071539 a=4187322124 w=2059
```

```
amatsu NCM-working-group(:|✓) % xxd -g 1 ./tcpka-v6.mdns
```

```

root@natsu NCM-working-group[~]:# python3 parse-rr-bin.py ./tcpka-v6.mdns 23
idx  name                                     type class ttl      rd_len
-----
0  dut-x-ssh-tcp.local.                     TXT 0x0001 4500 0x0001
1  _services-dns-sd-udp.local.               PTR 0x0001 4500 0x0002
2  _ssh-tcp.local.                           PTR 0x0001 4500 0x0002
3  dut-x-device-info-tcp.local.              TXT 0x0001 4500 0x0023
4  dut-x-sftp-ssh-tcp.local.                 TXT 0x0001 4500 0x0001
5  _services-dns-sd-udp.local.               PTR 0x0001 4500 0x0002
6  _sftp-ssh-tcp.local.                      PTR 0x0001 4500 0x0002
7  dut-x-rfb-tcp.local.                      TXT 0x0001 4500 0x0001
8  _services-dns-sd-udp.local.               PTR 0x0001 4500 0x0002
9  _rfb-tcp.local.                           PTR 0x0001 4500 0x0002
10 _kerberos.dut-newname.local.              TXT 0x0001 4500 0x0033
11 185EE7D6-A96A-442E-BE08-EC1AE11661F6._remotepairing-tcp.local. TXT 0x0001 4500 0x0059
12 _services-dns-sd-udp.local.               PTR 0x0001 4500 0x0002
13 _remotepairing-tcp.local.                 PTR 0x0001 4500 0x0002
14 dut-x-ssh-tcp.local.                     SRV 0x0001 4500 0x0008      SRV prio 0, weight 0, port 22, target dut-newname.local.
15 dut-x-sftp-ssh-tcp.local.                SRV 0x0001 4500 0x0008      SRV prio 0, weight 0, port 22, target dut-newname.local.
16 dut-x-rfb-tcp.local.                     SRV 0x0001 4500 0x0008      SRV prio 0, weight 0, port 5900, target dut-newname.local.
17 185EE7D6-A96A-442E-BE08-EC1AE11661F6._remotepairing-tcp.local. SRV 0x0001 4500 0x0008      SRV prio 0, weight 0, port 49152, target dut-newname.local.
18 5EB2E440-5A85-4855-8587-230D3DFACD9C._rp-tunnel-tcp.local.   TXT 0x0001 4500 0x0001
19 _services-dns-sd-udp.local.               PTR 0x0001 4500 0x0002
20 _rp-tunnel-tcp.local.                     PTR 0x0001 4500 0x0002
21 5EB2E440-5A85-4855-8587-230D3DFACD9C._rp-tunnel-tcp.local.   SRV 0x0001 4500 0x0008      SRV prio 0, weight 0, port 53531, target dut-newname.local.
22 152124.keapalive.dns-sd-udp.local.        NULL 0x0001 4500 0x0000      ts=15 i=30 c=10 H=df12:3456:789a:1::2 D=df12:3456:789a:1::1 l=49208 r=3003 m=9C:76:0E:31:00:8F s=1309919172 a=2154494200 w=2053

```

