

TCP Urgent Pointer

IETF 74 – TCPM WG

March 23, 2009

David Borman

Wind River

TCP Urgent pointer

- **draft-gont-tcpm-urgent-data-01.txt**
- **Mailing list discussion**

The original problem

- **RFC 793, pg. 17**
 - “The urgent pointer points to the sequence number of the octet following the urgent data.”
- **RFC 793, pg. 56**
 - “If the urgent flag is set, then $SND.UP \leftarrow SND.NXT-1$ and set the urgent pointer in the outgoing segments.”
- **RFC 793, pg. 48**
 - “Note that data following the urgent pointer (non-urgent data) cannot be delivered to the user in the same buffer with preceeding urgent data unless the boundary is clearly marked for the user.”

The corrections

- **RFC 961, pg. 6**

“Urgent: Page 17 is wrong. The urgent pointer points to the last octet of urgent data (not to the first octet of non-urgent data).”

- **RFC 1122, pg. 84**

“4.2.2.4 Urgent Pointer: RFC-793 Section 3.1

The second sentence is in error: the urgent pointer points to the sequence number of the LAST octet (not LAST+1) in a sequence of urgent data. The description on page 56 (last sentence) is correct.”

The problem

- **Most implementations use the RFC 793, pg. 17 definition**
 - **4.4 BSD-Lite2**
 - FreeBSD, NetBSD, OpenBSD, BSD/OS
 - **Linux**
 - **Microsoft Windows**
 - Windows 95, Windows 2000 SP4, Windows 2008
 - **Cisco IOS**
 - versions 12.2(18)SXF7, 12.4(15)T7
- **See draft-gont-tcpm-urgent-data-01.txt for more details**

4.4BSD-Lite2: sys/netinet/tcp_input.c

```
* According to RFC961 (Assigned Protocols),  
* the urgent pointer points to the last octet  
* of urgent data. We continue, however,  
* to consider it to indicate the first octet  
* of data past the urgent section as the original  
* spec states (in one of two places).  
*/  
if (SEQ_GT(ti->ti_seq+ti->ti_urp, tp->rcv_up)) {  
    tp->rcv_up = ti->ti_seq + ti->ti_urp;  
    so->so_oobmark = so->so_rcv.sb_cc +  
        (tp->rcv_up - tp->rcv_nxt) - 1;  
}
```

4.4BSD-Lite2: sys/netinet/tcp_usrreq.c

```
* According to RFC961 (Assigned Protocols),  
* the urgent pointer points to the last octet  
* of urgent data. We continue, however,  
* to consider it to indicate the first octet  
* of data past the urgent section.  
* Otherwise, snd_up should be one lower.  
*/  
sbappend(&so->so_snd, m);  
tp->snd_up = tp->snd_una + so->so_snd.sb_cc;
```

Out Of Band (OOB) Data

- TCP doesn't support OOB Data
- The BSD TCP implementation pulled a byte out of the TCP data stream at the urgent pointer
 - This is WRONG
 - Use the `SO_OOBINLINE` socket option to disable this behavior

Recommendations

- **Take this on as a WG item**
- **Use draft-gont-tcpm-urgent-data-01.txt as a starting point**
- **Change the definition of the Urgent Pointer to match RFC 793, pg. 17, since that is what is in most implementations.**
- **This would update RFC 793, RFC 961 and RFC 1122**

WIND RIVER