BGP configuration weirdness on Foundry/Brocade

If you're all too familiar with Cisco, then you will - as well as I do - struggle accross some weirdnesses on the Foundry/Brocade routers every now and then.

Not too long ago I fought around with a BGP issue on the XMR 4000.
My problem was that the XMR would announce just about any IPv6 prefix to all BGP peers, despite the fact that I had a configuration in place, which should effectifely only announce my own prefixes.

My Cisco configuration, which used to work properly, looks as shown below.
For the sake of simplicity I stripped away some advanced settings for communities, route-maps and prefix filters to keep it short.

```
ip as-path access-list 5 permit $
ip as-path access-list 5 deny .*
!
router bgp xxxxx
 neighbor PEERGROUPv6 peer-group
 neighbor PEERGROUPv6 description Some IPv6 Peers
 !
 neighbor SOMEIPV6PEER remote-as nnnnn
 neighbor SOMEIPV6PEER peer-group PEERGROUPv6
 !
 address-family ipv6
  neighbor PEERGROUPv6 soft-reconfiguration inbound
  neighbor PEERGROUPv6 filter-list 5 out
  !
  neighbor SOMEIPV6PEER activate
```

The basic idea of this setup is to have all peers share the same subset of settings through the peer-group, in this case the filter-list, which should only permit my own AS.

On the cisco, this has this exact effect, causing only my prefixes to be announced to the peers:

```
#show bgp ipv6 uni neighbors MASKEDIPV6PEER advertised-routes
BGP table version is 6144418, local router ID is MASKEDROUTER
Status codes: s suppressed, d damped, h history, * valid,  best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*iMASKEDIPV6PREFIX/32   MASKEDROUTER
                                  1    100      0 i
```

Doing the same configuration 1:1 on the Brocade XMR didn't work however. I always ended up having ALL IPv6 routes announced to my peers, no matter of their origin:

```
#show ipv6 bgp neigh MASKEDIPV6PEER advertised-routes
      There are 8366 routes advertised to neighbor MASKEDIPV6PEER
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST E:EBGP I:IBGP L:LOCAL
      Prefix            Next Hop        Metric     LocPrf     Weight Status
1     2001::/32         MASKEDIPV6PEER
                                990       400       0      BI
       AS_PATH: 6730 6695 12859
2     2001::/32         MASKEDIPV6PEER       301       0      E
       AS_PATH: 6939
3     2001::/32         MASKEDIPV6PEER
                                999       100       0      E
       AS_PATH: 13030 12859
4     2001:200::/32     MASKEDIPV6PEER
                                990       400       0      BI
       AS_PATH: 6730 6939 2500
5     2001:200::/32     MASKEDIPV6PEER 980       300       0      E
       AS_PATH: 6939 2500
6     2001:200::/32     MASKEDIPV6PEER
                                999       100       0      E
       AS_PATH: 13030 2500
```

If course this is not what I intended, as only my local prefixes should be announced.
My initial Brocade config looked very similar to the Cisco config:

```
ip as-path access-list 5 seq 5 permit $
ip as-path access-list 5 seq 25 deny .*
!
router bgp
 neighbor PEERGROUPv6 peer-group
 neighbor PEERGROUPv6 description "Some IPv6 Peers"
 !
 neighbor SOMEIPV6PEER remote-as nnnnn
 neighbor SOMEIPV6PEER peer-group PEERGROUPv6
 !
 address-family ipv6 unicast
  neighbor PEERGROUPv6 activate
  neighbor PEERGROUPv6 soft-reconfiguration inbound
  neighbor PEERGROUPv6 filter-list 5 out
  !
  neighbor SOMEIPV6PEER activate
```

The main difference between Cisco and Brocade, despite from some syntactical differences, was the need to have the "neighbor PEERGROUPv6 activate" statement in place.
So if all routes are announced to all peers, a misconfiguration within the peer-group would be most likely the case. I checked the peer-group configuration as follows:

```
#show ip bgp peer-group
1   BGP peer-group is PEERGROUPv6
    Description: Some IPv6 Peers
      NextHopSelf: no
      SoftInboundReconfiguration: yes
     Address family : IPV4 Unicast
    Route Filter Policies:
      Filter-list: (out) 5
     Address family : IPV4 Multicast
     Address family : IPV6 Unicast
      Filter-list: (out) 5
      Prefix-list: (in) ipv6-prefix-in  (out) ipv6-prefix-out
      Route-map: (in) IXin  (out) IXout
     Address family : IPV6 Multicast
    Members:
      IP Address: MASKEDIPV6PEER, AS: MASKEDAS
      IP Address: MASKEDIPV6PEER, AS: MASKEDAS
      IP Address: MASKEDIPV6PEER, AS: MASKEDAS
```

This looked ok to me. It didn't make any sense to me at all.

Well, I read the docs, but I didn't find any reasonable clue on this. Also friend Google dind't help.

After some time I had the idea to check what happens, if I add another "neighbor PEERGROUPv6 filter-list 5 out" command to the "address-family ipv4 unicast", see below:

```
ip as-path access-list 5 seq 5 permit $
ip as-path access-list 5 seq 25 deny .*
!
router bgp
 neighbor PEERGROUPv6 peer-group
 neighbor PEERGROUPv6 description "Some IPv6 Peers"
 !
 neighbor SOMEIPV6PEER remote-as nnnnn
 neighbor SOMEIPV6PEER peer-group PEERGROUPv6
 !
 address-family ipv4 unicast
  neighbor PEERGROUPv6 filter-list 5 out
 !
 address-family ipv6 unicast
  neighbor PEERGROUPv6 activate
  neighbor PEERGROUPv6 soft-reconfiguration inbound
  neighbor PEERGROUPv6 filter-list 5 out
  !
  neighbor SOMEIPV6PEER activate
```

Now I would need to reset the BGP peer to see if the change had an effect:

#clear ipv6 bgp neighbor SOMEIPV6PEER soft-outbound

I couldn't believe it when I saw the result:

```
#show ipv6 bgp neigh MASKEDIPV6PEER advertised-routes
      There are 1 routes advertised to neighbor 2001:7f8:24::aa
Status A:AGGREGATE B:BEST b:NOT-INSTALLED-BEST E:EBGP I:IBGP L:LOCAL
      Prefix            Next Hop        Metric     LocPrf     Weight Status
1     MASKEDIPV6PREFIX::/32     MASKEDROUTER         0      BE
       AS_PATH: MASKEDAS
```

To see if I had done something wrong, I reverted all changes and tried again.
Result: Again all routes were announced.

Having reapplied the "IPv4 filter-list" again caused only required prefixes to be announced -- exactly the way how I intended it.

Now, this doens't make any sense at all and feels like a software bug within the implementation.
Even before adding this non-obvious configuration command, the output of "show ip bgp peer-group" clearly stated which filters and prefix lists the peer-group used.
The output didn't change at all after applying the changed configuration.

At least my BGP announcements are correct now. The case still needs to be resolved with the Brocade tech guys.