ASG v7 Cluster

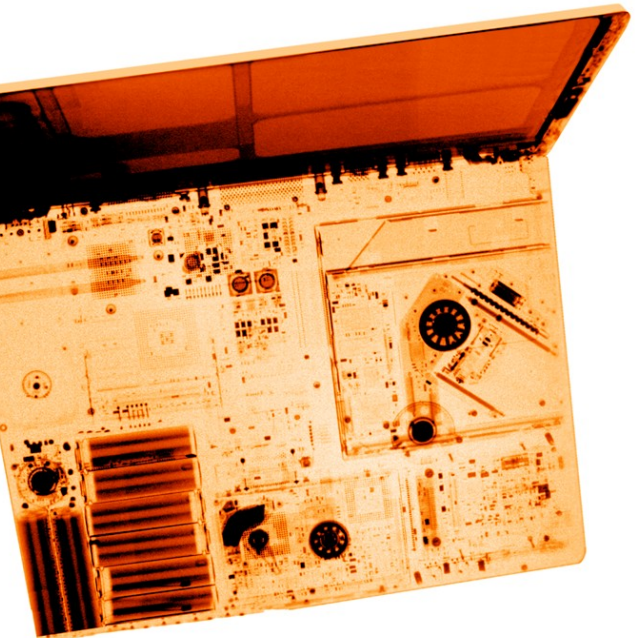

Ulrich Weber

Requirements

- High Availability / Failover
- No external Load Balancer
- Whole cluster as one unit with one IP
- Distribute high CPU intensive tasks

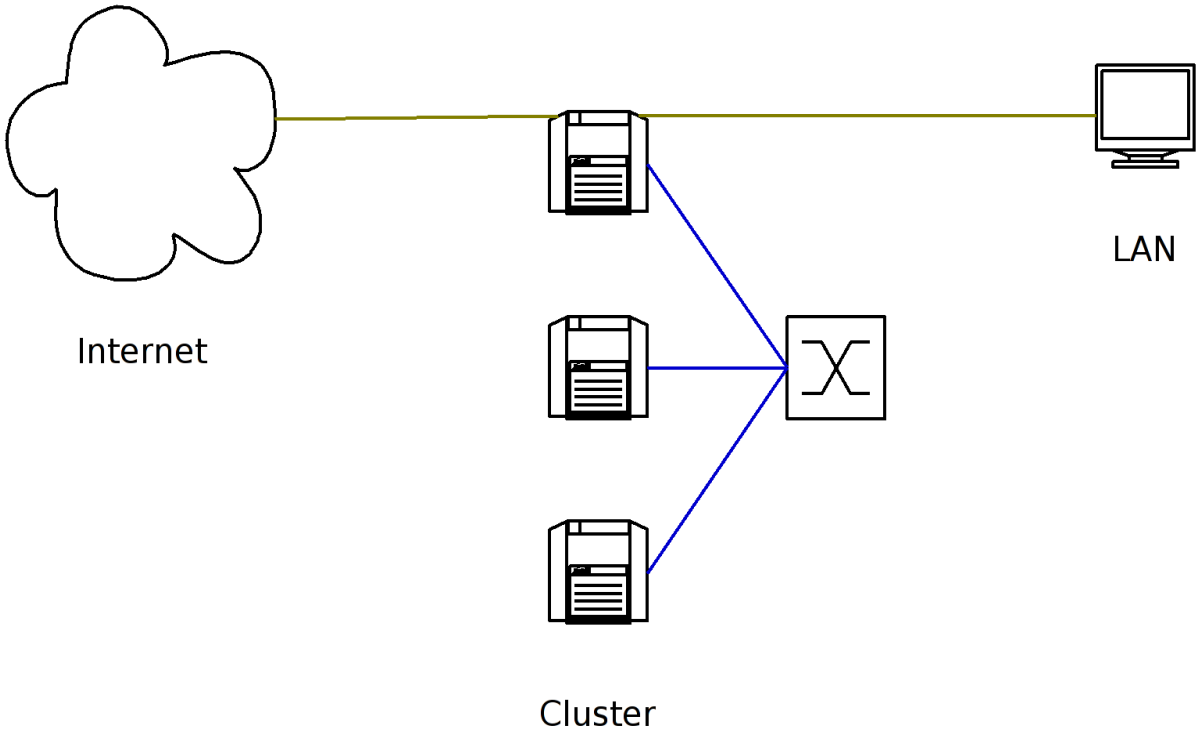

Distribution

- IPS - packet analyzing
- IPSec - encryption/decryption
- Proxies
  - HTTP
  - SMTP
  - POP3
  - FTP

Distribution Algorithms

- Hash (Source IP)
- Round Robin

WARP basics

- Grab / Reinject packets within chains
- Based on ip conntrack
  - Extended *ip_conntrack* struct with *u_int8_t node_id;*
  - Extended *sk_buf* struct with *__u32 id;*
  - Works with helpers

- Mute interfaces within *netif_receive_skb* and *dev_queue_xmit* on Slave nodes
- Jumbo Frame Support
- Unique port range for each cluster node

```
 enum nf_ip_hook_priorities {
        NF_IP_PRI_FIRST = INT_MIN,
+       NF_IP_PRI_IPS_OUT = INT_MIN + 1, /* ips_table (LOCAL_OUT) */
        NF_IP_PRI_CONNTRACK_DEFRAG = -400,
        NF_IP_PRI_RAW = -300,
        NF_IP_PRI_SELINUX_FIRST = -225,
@@ -64,12 +65,20 @@
        NF_IP_PRI_BRIDGE_SABOTAGE_LOCAL_OUT = -50,
        NF_IP_PRI_FILTER = 0,
        NF_IP_PRI_NAT_SRC = 100,
-       NF_IP_PRI_IPS = 200,                    /* ips table (low priority)*/
-       NF_IP_PRI_IPS_OUT = INT_MIN + 1, /* ips_table (LOCAL_OUT) */
+       NF_IP_PRI_IPS_FORW = 200, /* ips_table (FORWARD) */
        NF_IP_PRI_SELINUX_LAST = 225,
+#ifdef CONFIG_ASG_CLUSTER
+       NF_IP_PRI_CONNTRACK_HELPER = INT_MAX - 5,
+       NF_IP_PRI_NAT_SEQ_ADJUST = INT_MAX - 4,
+       NF_IP_PRI_CONNTRACK_CONFIRM = INT_MAX - 3,
+       NF_IP_PRI_CLUSTER = INT_MAX - 2,
+       NF_IP_PRI_IPS_IN = INT_MAX - 1, /* ips_table (LOCAL_IN) */
+#else
+       NF_IP_PRI_IPS_IN = 200, /* ips_table (LOCAL_IN) */
        NF_IP_PRI_CONNTRACK_HELPER = INT_MAX - 2,
        NF_IP_PRI_NAT_SEQ_ADJUST = INT_MAX - 1,
        NF_IP_PRI_CONNTRACK_CONFIRM = INT_MAX,
+#endif
        NF_IP_PRI_LAST = INT_MAX,
 };
```

WARP Header

```
struct warphdr {
        u16                     len;
        u8                      flags;
        u8                      hook;
        u32                     hook_priority;
        u32                     id;
        unsigned char           dev[IFNAMSIZ];
        struct skb_timeval      tstamp;
#ifdef CONFIG_NET_CLS_ACT
        unsigned char           input_dev[IFNAMSIZ];
#endif
        u8                      cb[48];
        u32                     priority;
        u32                     nfmark;
        u32                     ipsec_seq;
        u32                     redirect_addr;
        u16                     redirect_port;
        u32                     mac_len;
        u16                     h_off;
        u8                      local_df;
};
```

Demonstration

- IPSec Distribution

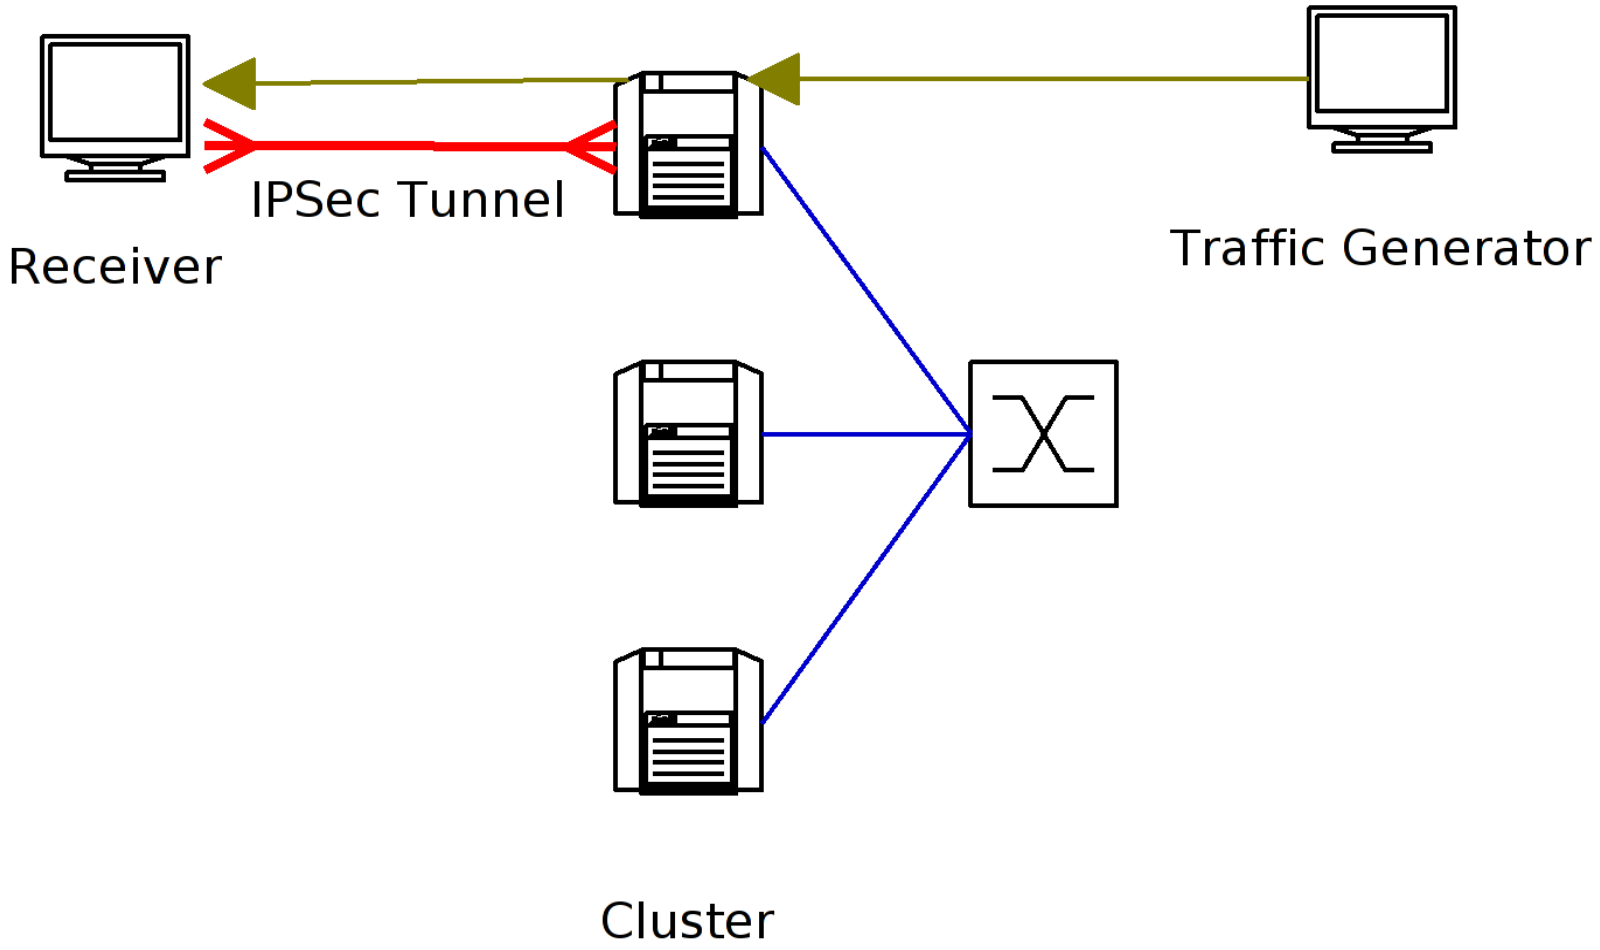

Fine

Questions ?