Implications of IPv6 Addressing on Security Operations

(draft-gont-opsec-ipv6-addressing)

F. Gont, SI6 Networks
G. Gont, SI6 Networks

OPSEC WG. IETF 116
March 25rd-31st, 2023

Introduction

Motivation

- Differences in IPv6 vs. IPv4 addressing have concrete implications on security operations
- These might be non-obvious outside of IPv6-savvy circles, e.g.,
  - Cloud operations groups
  - Security operations groups
- Such groups continue applying IPv4 practices -> fail!
- This document has been motivated by conversations with such groups

What do we mean by IPv6 addressing “differences”?

- IPv6 addresses have an associated:
  - address scope: global, link-local, etc.
  - stability property: stable vs. temporary
  - intended usage: outgoing vs. incoming communications
- IPv6 nodes typically use multiple addresses simultaneously
- IPv6 addresses may be configured via different mechanisms:
  - Configuration mechanism may have implications on the address properties
- IPv6 users typically control a large IPv6 address block (e.g. a /64)

What is behind an IPv6 prefix?

- Multiple addresses may map to a single host
  - Host typically configure multiple addreses
  - Addresses typically selected from /64
  - But a user might control a larger address block (e.g. a whole /48)
- A single IPv6 address may map to multiple hosts
  - NAT-PT for IPv6 not uncommon
  - Kubernetes typically do IPv6 ULAs + NAT
- All these aspects are key when doing IPv6 security operations

IPv6 security operations

- Enforcing Access Control Lists (ACLs):
  - Allow-lists:
    - Meant to allow access from a single system or group of systems
  - Block-lists:
    - Meant to block access from a single system or group of systems
- Network activity correlation
  - Analyze relationship between different network activities

IPv6 Security Operations

Challenges

ACLs: Allow-lists

- Use of temporary addresses (RFC8981) means:
  - Addresses change on a regular basis
  - Addresses from multiple hosts may be intermingled in the same /64
- But...What should we “allow”?
- If specifying /128s, the ACLs might fail

ACLs: Block-lists

- Quite often, these are dynamically introduced, e.g.
  - SIEM/IPS
  - fail2ban
  - IP reputation services (e.g., abuseipdb.com)
- But...what should we “block”?
- If blocking /128s, a skilled attacker might:
  - Intentionally exhaust the number of entries in your block-list
  - Circumvent the block-list (i.e., use *throw-away* IPv6 addresses)

Network Activity Correlation

- Non-trivial exercise:
  - Multiple systems might be behind a /128, or,
  - A single system might jump around within a /48, or,
  - Anything in between

IPv6 Security Operations

Possible Advice

ACLs: Allow-lists

- Employ stable addresses (only):
  - Use:
    - manual configuration, or,
    - DHCPv6, or,
    - SLAAC & disable temporary addresses (e.g. via group policies)
  - Specify allow-lists as /128s
- Embrace temporary addresses usage:
  - Segregate systems into different subnets
  - Specify allow-lists as,e.g., /64s

ACLs: Block-lists

- If block-lists are dynamically-generated:
  - May need to dynamically aggregate ACLs
  - Possibly adjust the ACL lifetime based on the aggregation level

ACLs: Block-lists (II)

- One possible implementation for dynamic block-lists:

LEVEL	PREF_LEN	AGGR_THRES	ACL_LIFETIME
1	/128	10	1 hour
2	/64	10	1 hour
3	/56	10	30 min
4	/48	N/A	15 min

"Where possible, agregate at least $\texttt{AGGR\_THRES}_{N}$ $\texttt{LEVEL}_{N}$ ACLs into a single $\texttt{LEVEL}_{(N+1)}$ ACL. Remove this new ACL after $\texttt{ACL\_LIFETIME}_{(N+1)}$"

Network Activity Correlation

- Tools should readily allow activity correlation on a per-prefix basis

Moving forward

Moving forward

- Comments?
- Is there interest in:
  - working on this topic?
  - adopting this document as a wg item?