Background Concepts for Network Security Readings

Prerequisite concepts for assigned papers

Roadmap

Part 1

Infrastructure & Wireless

- CDN Architecture & DoS
  → *CDN Judo (NDSS 2020)*
- Inter-Domain Routing & Path Validation
  → *Symphony (NDSS 2024)*
- Wi-Fi Security & 802.11
  → *Framing Frames (USENIX Sec 2023)*

Part 2

Detection, Inspection & Censorship

- Intrusion Detection & Explainable AI
  → *xNIDS (USENIX Sec 2023)*
- Stateful DPI & TCP Evasion
  → *SymTCP (NDSS 2020)*
- Internet Censorship & Circumvention
  → *DeResistor (USENIX Sec 2023)*

Part 1: Infrastructure & Wireless

CDN architecture, inter-domain routing, and Wi-Fi security

Content Delivery Networks

How CDNs work and how they can be abused

THE UNIVERSITY OF MELBOURNE

```
nikita@wirelessprv-10-193-126-83 ~/work/courses/cs461-ece422/fa25/slides $ ping www.unimelb.edu.au
PING uom-web02.uom-7329.saas.squiz.cloud (2.58.104.10): 56 data bytes
64 bytes from 2.58.104.10: icmp_seq=0 ttl=52 time=13.050 ms
64 bytes from 2.58.104.10: icmp_seq=1 ttl=52 time=44.688 ms
64 bytes from 2.58.104.10: icmp_seq=2 ttl=52 time=88.060 ms
64 bytes from 2.58.104.10: icmp_seq=3 ttl=52 time=13.560 ms
64 bytes from 2.58.104.10: icmp_seq=4 ttl=52 time=14.956 ms
64 bytes from 2.58.104.10: icmp_seq=5 ttl=52 time=11.291 ms
64 bytes from 2.58.104.10: icmp_seq=6 ttl=52 time=11.434 ms
64 bytes from 2.58.104.10: icmp_seq=7 ttl=52 time=12.545 ms
64 bytes from 2.58.104.10: icmp_seq=8 ttl=52 time=16.360 ms
64 bytes from 2.58.104.10: icmp_seq=9 ttl=52 time=11.637 ms
^C
```

Accept your offer and enrol

→ Get started

CDN Goals

- **Geographic proximity**: data centers around the world
- **DoS protection**: big pipes, robot detection
- **Caching**: reduce origin load

aws

Akamai

CLOUDFLARE

fastly

CDN operation

Request Path

- Client uses DNS/anycast to locate nearest CDN point of presence
- CDN PoP checks cache
- Upon miss, forwards to origin over **separate** connection
  - Different versions of HTTP, TLS
  - Request is modified

DoS / DDoS Attack Models

Volumetric Attacks

- Flood target with massive traffic volume
- UDP floods, DNS amplification, NTP reflection
- Amplification: small request → large response reflected to victim
- Measured in Gbps or Mpps
- CDNs are designed to absorb these

Application-Layer Attacks

- HTTP flood: legitimate-looking requests that exhaust server resources
- Slowloris: hold connections open slowly
- Cache-busting: unique URLs bypass CDN cache
- Harder to distinguish from real traffic
- CDN Judo exploits CDN's own mechanisms for amplification

Inter-Domain Routing

BGP, path validation, and securing the routing infrastructure

Autonomous Systems

- The Internet is composed of tens of thousands of *autonomous systems* (ASs)
- Each AS *controls* its own network
- ASes collaborate to enable global delivery while fulfilling individual goals
  - Maintaining quality of service
  - Relationship management
  - Commercial considerations
  - …

Inter- vs Intra-domain Routing

- Intra-domain routing: routing within a single autonomous system
  - Can use shortest (weighted) path algorithms
- Inter-domain routing: routing between autonomous systems
  - Supported by the Border Gateway Protocol (BGP)
  - BGP distributes global knowledge of routes while allowing each AS to apply its own *policies*

BGP Route Announcements

- BGP lets border routers *announce* routes to each other
- Route:
  - IP Prefix: 1.2.3/24 — 1.2.3.xx, 1.2/16 — 1.2.x.x, etc.
  - AS Path: A list of AS (numbers) that are used to deliver to prefix
- Example:
  - 140.141.128.0/17 34854 (Staclar) 6939 (HE) 12119 (i3)
- Route announcement = offer to deliver traffic to prefix

BGP Lifecycle

- Border router receives route announcement from neighbors
- Uses route to modify *forwarding tables* of internal routers
  - A series of (prefix, next hop) pairs
  - Routers use *most specific* prefix: 1.2.3/24 beats 1.2/16
- As prefix arrives to other border routers, it is announced to other neighbors
  - With current AS added to path
- Rinse, repeat

BGP Example

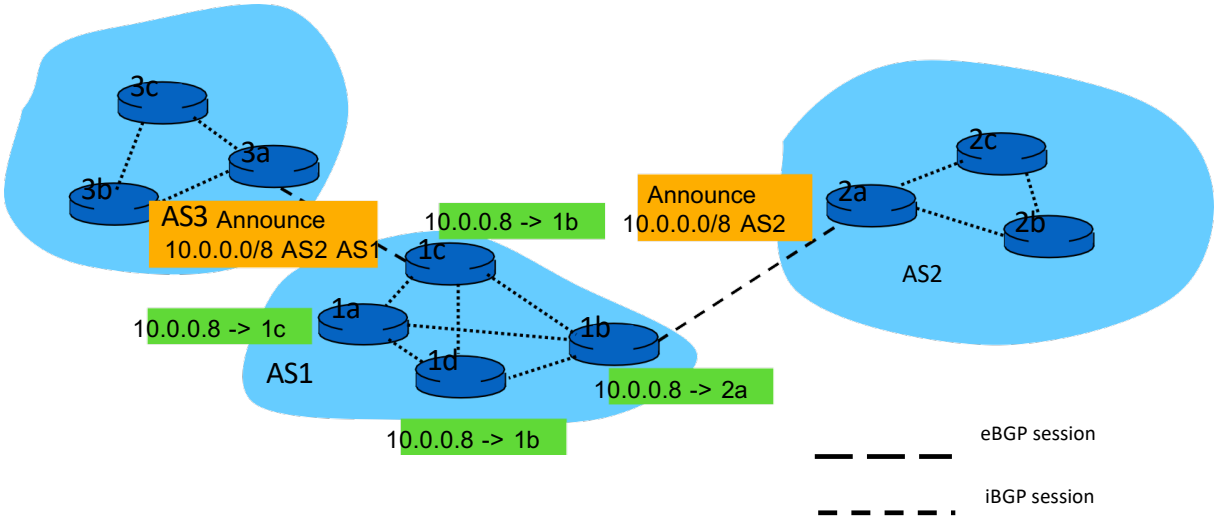
3c
3a
3b
AS3 Announce
10.0.0.0/8 AS2 AS1
10.0.0.8 -> 1b
Announce
10.0.0.0/8 AS2
2c
2a
2b
AS2
1c
1a
1b
1d
10.0.0.8 -> 1c
AS1
10.0.0.8 -> 2a
10.0.0.8 -> 1b
eBGP session
iBGP session

BGP Lifecycle

- Border router receives route announcement from neighbors
- Uses route to modify *forwarding tables* of internal routers
  - A series of (prefix, next hop) pairs
  - Routers use *most specific* prefix: 1.2.3/24 beats 1.2/16
- As prefix arrives to other border routers, it is announced to other neighbors
  - With current AS added to path
- Rinse, repeat

Import policy

Routing policy

Export policy

BGP Hijacking

BGP attacks hijack Telegram traffic in Iran

With so many users in Iran, it's unsurprising that potentially state-sponsored groups would want an access point into the banned app.

Pakistan's Accidental YouTube Re-Routing Exposes Trust Flaw in Net

A Pakistan ISP that was ordered to censor YouTube accidentally managed to take down the video site arou
the world for several hours Sunday. The Pakistani government ordered ISPs to censor YouTube to prevent
Pakistanis from seeing a trailer to an anti-Islamic film by Dutch politician Geert Wilders. YouTube has since
removed the clip for violating its terms of service, but a screenshot [...]

Popular Destinations rerouted to Russia

Posted by Andree Toonk - December 12, 2017 - Hijack - No Comments

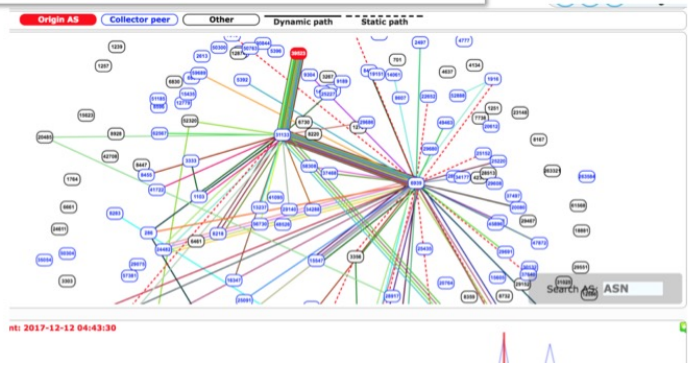

Path-Aware Networking

- Traditional Internet: source has no control over path; routing is hop-by-hop
- Source routing: sender specifies the path the packet should take
- SCION architecture: divides Internet into Isolation Domains (ISDs) with cryptographic path control
- Path validation: verifying that packets actually follow the declared path
- Challenge: per-packet verification is expensive at line rate (100+ Gbps)
- Cryptographic hop fields: each AS embeds a MAC so downstream ASes can verify

Wi-Fi Security

802.11 fundamentals, encryption, and frame handling

802.11 Frame Types

- Management frames: beacons, probe requests/responses, authentication, association, deauthentication
- Control frames: RTS/CTS, ACK — coordinate access to the wireless medium
- Data frames: carry actual payload; can be individually or group-addressed
- Frame aggregation (A-MSDU, A-MPDU): bundle multiple frames for throughput
- A-MSDU: multiple MSDUs in one MPDU — aggregation happens above encryption
- A-MPDU: multiple MPDUs in one transmission — each independently encrypted

WPA2 / WPA3 Key Hierarchy

WPA2 (RSN)

- 4-way handshake derives Pairwise Transient Key (PTK)
- PTK protects unicast traffic between client and AP
- Group Temporal Key (GTK) protects broadcast/multicast
- Vulnerabilities: KRACK attack showed nonce reuse in handshake
- TKIP deprecated; CCMP (AES) is standard

WPA3 (SAE)

- Simultaneous Authentication of Equals (SAE) replaces PSK exchange
- Forward secrecy: compromised password doesn't expose past traffic
- Protected Management Frames (PMF) now mandatory
- IGTK: Integrity Group Temporal Key for broadcast management frames
- Dragonfly handshake resistant to offline dictionary attacks

Power-Save Mode & Transmit Queues

- Clients can enter power-save mode to conserve battery — AP buffers frames
- AP maintains per-client transmit queues for buffered frames
- Client signals wake-up; AP delivers buffered frames
- Key issue: what security context are buffered frames encrypted with?
- If client re-associates with new keys, what happens to frames queued under old keys?
- Authentication/association state and encryption state are managed separately in 802.11
- This gap between state machines is central to the Framing Frames attack

Part 2: Detection, Inspection & Censorship

Intrusion detection, deep packet inspection, and Internet censorship

Network Intrusion Detection

Architectures, machine learning, and explainability

NIDS Architectures

Signature-Based

- Match traffic against database of known attack patterns
- Low false positive rate for known attacks
- Cannot detect zero-day or novel attacks
- Requires constant signature updates
- Example: Snort, Suricata rule sets

Anomaly-Based

- Build model of 'normal' traffic; flag deviations
- Can detect novel attacks
- Higher false positive rate
- Requires training on representative baseline
- ML-based NIDS fall in this category

Deployment: Inline vs. Passive

Inline (IPS)

- Sits in the network path — all traffic passes through
- Can actively block or drop malicious packets
- Adds latency; failure = network outage
- Must make decisions in real time
- False positives directly block legitimate traffic

Passive (IDS)

- Monitors a copy of traffic (span port, TAP)
- Alerts on suspicious activity but doesn't block
- No added latency or risk of outage
- More time for complex analysis
- Response is manual or via integration with firewalls

Deep Learning for Traffic Classification

- Input representations: raw packet bytes, packet headers, flow-level statistics (duration, byte count)
- CNNs: treat packet payloads or header sequences as spatial data
- RNNs/LSTMs: model temporal patterns in packet sequences within a flow
- Feature engineering vs. end-to-end learning: hand-crafted features vs. raw input
- Challenges: encrypted traffic limits payload inspection; models trained on outdated datasets
- Alert fatigue: even 0.1% false positive rate = thousands of false alarms per day at scale

Explainable AI for Security

- Black-box DL models are hard to trust: why did the model flag this flow?
- SHAP (SHapley Additive exPlanations): assigns contribution scores to each input feature
- Local explanations: why this specific prediction; global: what features matter overall
- For NIDS: explanations should map to actionable network defense (block IP, port, protocol)
- Challenge: translating feature attributions into concrete firewall rules or response actions
- Interpretability vs. accuracy tradeoff: simpler models are more explainable but less accurate

Stateful Deep Packet Inspection

TCP internals, middlebox behavior, and evasion techniques

TCP State Machine Review

- Three-way handshake: SYN → SYN-ACK → ACK (connection established)
- Data transfer: sequence numbers, acknowledgments, windowing
- Connection teardown: FIN/ACK exchange or RST (abrupt reset)
- Retransmissions: lost segments are re-sent; receiver must handle duplicates
- Overlapping segments: what happens when retransmitted data differs from original?
- OS implementations differ in how they resolve ambiguities (first vs. last wins)

How Stateful DPI Works

- Middlebox reconstructs TCP streams by tracking connection state for every flow
- Reassembles payloads in order to inspect application-layer content (HTTP, DNS, etc.)
- Must handle out-of-order packets, retransmissions, and fragmentation
- Maintains per-connection state: sequence numbers, reassembly buffers
- Used by: firewalls, nation-state censors, corporate security appliances, ISP lawful intercept
- Resource intensive: must track millions of concurrent connections

Evasion vs. Insertion Attacks

Evasion

- Attacker sends packets that the IDS/DPI accepts but the end host drops
- IDS thinks it saw the full stream but missed the real payload
- Example: packet with low TTL — reaches IDS but not the server
- IDS reassembles stream incorrectly

Insertion

- Attacker sends packets that the IDS drops but the end host accepts
- IDS sees extra data that dilutes or obscures the real payload
- Example: overlapping TCP segment with different data
- IDS reassembles differently than the end host

Internet Censorship

How censors operate and how circumvention tools fight back

Censorship Techniques

- DNS poisoning: return false DNS responses for blocked domains
- IP blocking: blackhole traffic to known circumvention server IPs
- SNI-based blocking: inspect TLS Client Hello for the server name
- Active probing: censor connects to suspected circumvention servers to fingerprint them
- DPI-based blocking: inspect packet payloads for protocol signatures (e.g., Tor handshake)
- Throttling: degrade performance rather than fully block (harder to detect as censorship)

Circumvention Approaches

Tunneling & Proxying

- VPNs: encrypt all traffic through an exit point outside the censored region
- Tor: onion routing through multiple relays for anonymity + circumvention
- Domain fronting: use a CDN's domain in SNI while routing to a different backend
- Pluggable transports: make Tor traffic look like other protocols (e.g., obfs4, meek)

Detection Resistance

- Traffic must be indistinguishable from 'normal' allowed traffic
- Protocol mimicry: look like HTTPS, WebSocket, or video streaming
- Randomization: avoid fixed byte patterns that DPI can match
- Probe resistance: server must not reveal its nature when probed by the censor
- Cat-and-mouse: each improvement triggers new detection methods

Active Probing in Detail

- Censor identifies suspicious connection (e.g., unusual TLS fingerprint)
- Censor replays or modifies the client's initial bytes to the server ('probe')
- If the server responds like a circumvention tool, it gets blocked
- Probe types: exact replay, random payload, protocol-specific handshake
- Defense: circumvention servers must be 'probe-resistant' — refuse to respond to invalid clients
- Challenge: distinguishing legitimate clients from censor probes without breaking usability

Threat Modeling & Attacker Models

- Every paper assumes a specific attacker: define capabilities, position, and goals
- On-path vs. off-path: can the attacker see/modify traffic in transit?
- Local vs. remote: is the attacker on the same network or across the Internet?
- Passive vs. active: observing only, or injecting/modifying packets?
- Resource assumptions: nation-state (unlimited) vs. individual (limited)
- Always ask: what changes if we weaken or strengthen the threat model?

Research Methodology & Ethics

- Responsible disclosure: authors notify vendors before publication
- Ethical considerations: censorship evasion research has dual-use implications
- Evaluation methodology: real-world testing vs. lab simulation vs. formal analysis
- Reproducibility: are the tools/code released? Can results be independently verified?
- Measurement ethics: probing live censorship systems may put users at risk
- Arms race framing: many of these papers represent one move in an ongoing back-and-forth