Internet Security Introduction

ISP Workshops

Last updated 4th August 2018

Acknowledgements

- This material originated from the Cisco ISP/IXP Workshop Programme developed by Philip Smith & Barry Greene
  - These slides were originally developed by Dean Pemberton

- Bug fixes and improvements are welcomed
  - Please email *workshop (at) bgp4all.com*

Philip Smith

Introduction to Internet Infrastructure Security

- Introduction to the main network security issues that infrastructure operators need to be aware of.
- This includes discussion on packet flooding, Internet worms, DDOS attacks and Botnets

Why Security

Why do we need security on the Internet?

Why Security?

- The Internet was initially designed for connectivity
  - Trust is assumed, no security
  - Security protocols added on top of the TCP/IP
- Fundamental aspects of information must be protected
  - Confidential data
  - Employee information
  - Business models
  - Protect identity and resources
- The Internet has become fundamental to our daily activities (business, work, and personal)

Internet Evolution

LAN connectivity

Application-specific
More online content

Application/data
hosted in the "cloud"

- Different ways to handle security as the Internet evolves

ACRONYM/TERM OVERLOAD

- C I A
  - Confidentiality
  - Integrity
  - Availability
- Access Control
  - Authentication
  - Authorisation
  - Accountability
- Risk
- Threat
- Vulnerability
- Impact

Goals of Information Security

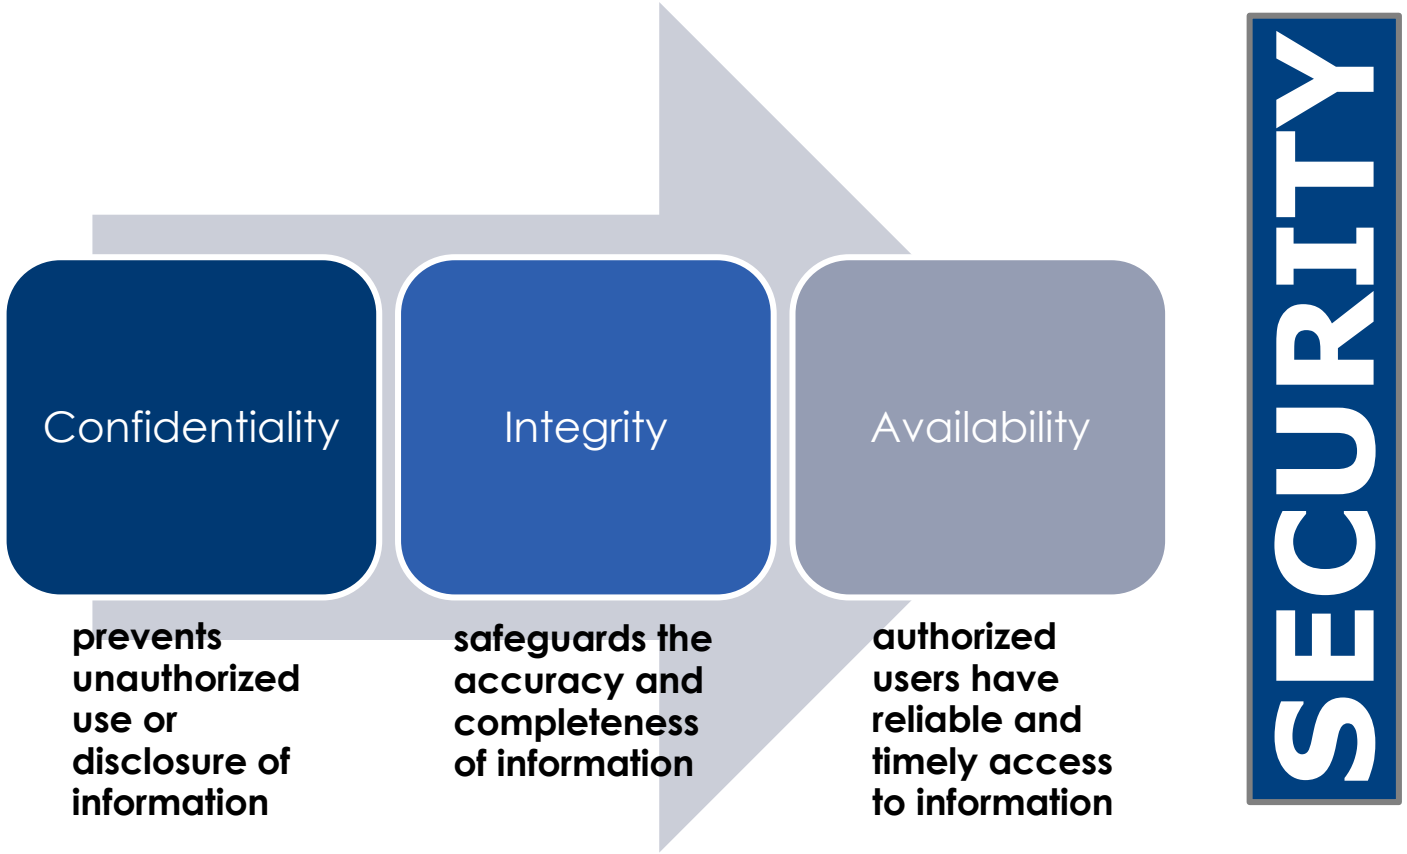

Access Control

- The ability to permit or deny the use of an object by a subject.

- It provides 3 essential services:
  - Authentication (identification of a user)
  - Authorisation (who is allowed to use a service)
  - Accountability (what did a user do)

Authentication

- A means to verify or prove a user's identity
- The term "user" may refer to:
  - Person
  - Application or process
  - Machine or device
- Identification comes before authentication
  - Provide username to establish user's identity
- To prove identity, a user must present either of the following:
  - What you know (passwords, passphrase, PIN)
  - What you have (token, smart cards, passcodes, RFID)
  - Who you are (biometrics such as fingerprints and iris scan, signature or voice)

Authentication – Examples of Tokens

eToken

RFID cards

Smart Cards

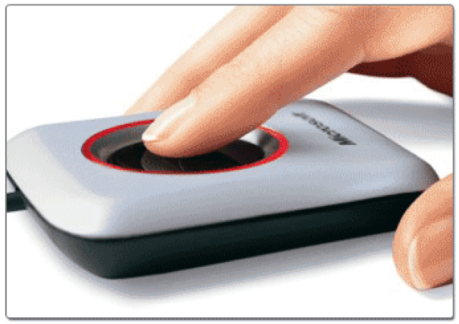

Fingerprint scanner

Authentication - Trusted Network

- Standard defensive-oriented technologies
  - Firewall – first line of defense
  - Intrusion Detection – second line of defense
- Build TRUST on top of the TCP/IP infrastructure
  - Strong authentication
    - Two-factor authentication
    - something you have + something you know
  - Public Key Infrastructure (PKI)

Strong Authentication

- An absolute requirement
- Two-factor authentication
  - Passwords (something you know)
  - Tokens (something you have)
- Examples:
  - Passwords
  - Tokens
  - Tickets
  - Restricted access
  - PINs
  - Biometrics
  - Certificates

Two-factor Authentication

- Requires a user to provide at least two authentication 'factors' to prove his identity
  - something you know
    - Username/userID and password
  - something you have
    - Token using a one-time password (OTP)
- The OTP is generated using a small electronic device in physical possession of the user
  - Different OTP generated each time and expires after some time
  - An alternative way is through applications installed on your mobile device
- Multi-factor authentication is also common

Authorisation

- Defines the user's rights and permissions on a system
- Typically done after user has been authenticated
- Grants a user access to a particular resource and what actions he is permitted to perform on that resource
- Access criteria based on the level of trust:
  - Roles
  - Groups
  - Location
  - Time
  - Transaction type

Authentication vs. Authorisation

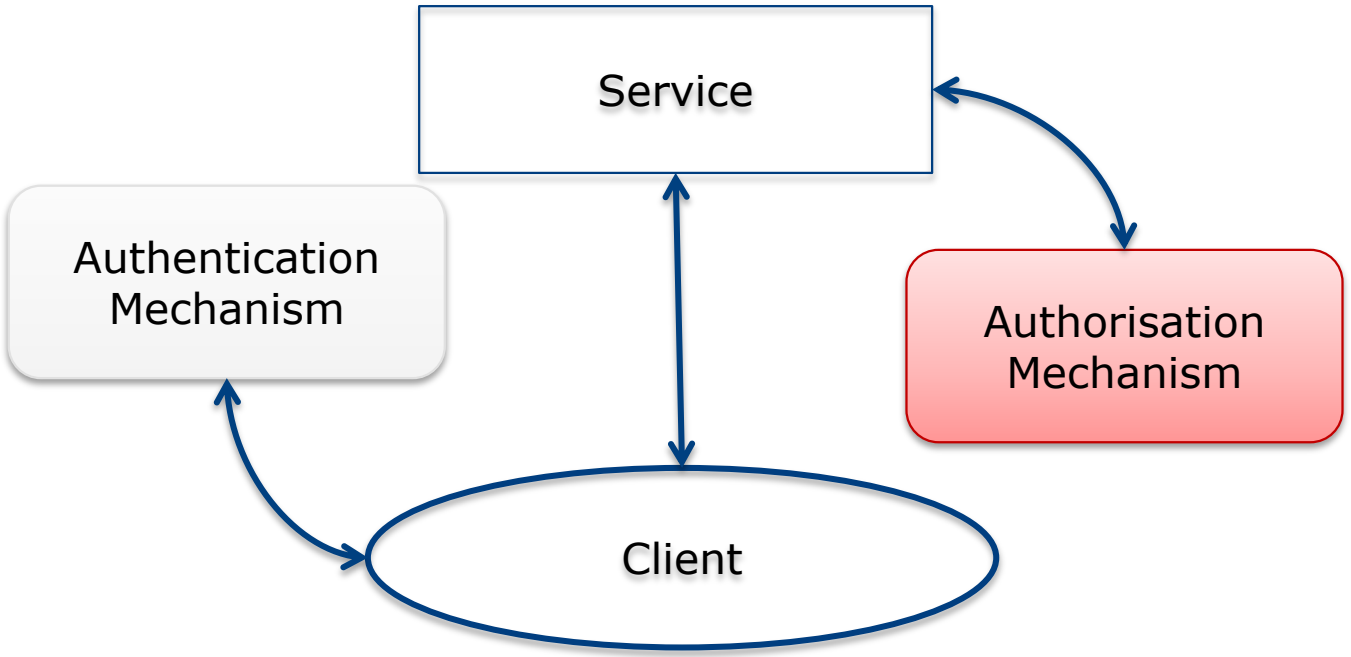

"Authentication simply identifies a party, Authorisation defines whether they can perform certain action" – RFC 3552

Authorisation Concepts

- Authorisation Creep
  - When users may possess unnecessarily high access privileges within an organization
- Default to Zero
  - Start with zero access and build on top of that
- Need to Know Principle
  - Least privilege; give access only to information that the user absolutely need
- Access Control Lists
  - List of users allowed to perform particular access to an object (read, write, execute, modify)

Authorisation – Single Sign On

- Property of access control where a user logs in only once and gains access to all authorized resources within a system.
- Benefits:
  - Ease of use
  - Reduces logon cycle (time spent re-entering passwords for the same identity)
- Common SSO technologies:
  - Kerberos, RADIUS
  - Smart card based
  - OTP Token
  - Shibboleth / SAML
  - OpenID
- Disadvantage: Single point of attack

Authorisation – Types of Access Control

- Centralized Access Control
  - Radius
  - TACACS+
  - Diameter
- Decentralized Access Control
  - Control of access by people who are closer to the resources
  - No method for consistent control

Accountability

- The security goal that generates the requirement for actions of an entity to be traced uniquely to that entity
  - Senders cannot deny sending information
  - Receivers cannot deny receiving it
  - Users cannot deny performing a certain action
- Supports nonrepudiation, deterrence, fault isolation, intrusion detection and prevention and after-action recovery and legal action

Source: NIST Risk Management Guide for Information Technology Systems

Integrity

- Security goal that generates the requirement for protection against either intentional or accidental attempts to violate data integrity
- Data integrity
  - The property that data has when it has not been altered in an unauthorized manner
- System integrity
  - The quality that a system has when it performs its intended function in an unimpaired manner, free from unauthorized manipulation

Source: NIST Risk Management Guide for Information Technology Systems

Risk, Threats, and Vulnerability

- Threat
  - Any circumstance or event with the potential to cause harm to a networked system
- Vulnerability
  - A weakness in security procedures, network design, or implementation that can be exploited to violate a corporate security policy
- Risk
  - The possibility that a particular vulnerability will be exploited

Threat

- “a motivated, capable adversary”
- Examples:
  - Human Threats
    - Intentional or unintentional
    - Malicious or benign
  - Natural Threats
    - Earthquakes, tornadoes, floods, landslides
  - Environmental Threats
    - Long-term power failure, pollution, liquid leakage

Vulnerability

- A weakness in security procedures, network design, or implementation that can be exploited to violate a corporate security policy
  - Software bugs
  - Configuration mistakes
  - Network design flaw
  - Lack of encryption
- Where to check for vulnerabilities?
- Exploit
  - Taking advantage of a vulnerability

Risk

- Likelihood that a vulnerability will be exploited
- Some questions:
  - How likely is it to happen?
  - What is the level of risk if we decide to do nothing?
  - Will it result in data loss?
  - What is the impact on the reputation of the company?
- Categories:
  - High, medium or low risk

Risk = Threat * Vulnerability * Impact

What Can Intruders Do?

- Eavesdrop - compromise routers, links, or DNS
- Send arbitrary messages (spoof IP headers and options)
- Replay recorded messages
- Modify messages in transit
- Write malicious code and trick people into running it
- Exploit bugs in software to ‘take over’ machines and use them as a base for future attacks

What are Security Goals?

- Controlling Data Access
- Controlling Network Access
- Protecting Information in Transit
- Ensuring Network Availability
- Preventing Intrusions
- Responding To Incidences

Goals are Determined by

- Services offered vs. security provided
  - Each service offers its own security risk
- Ease of use vs. security
  - Easiest system to use allows access to any user without password
- Cost of security vs. risk of loss
  - Cost to maintain

Goals must be communicated to all users, staff, managers, through a set of security rules called "security policy"

Causes of Security Related Issues

- Protocol error
  - No one gets it right the first time
- Software bugs
  - Is it a bug or feature ?
- Active attack
  - Target control/management plane
  - Target data plane
  - More probable than you think !
- Configuration mistakes
  - Most common form of problem

Why Worry About Security?

- How much you worry depends on risk assessment analysis
  - Risk analysis: the process of identifying security risks, determining their impact, and identifying areas requiring protection
- Must compare need to protect asset with implementation costs
- Define an effective security policy with incident handling procedures

Characteristics of a Good Policy

- Can it be implemented technically?
- Are you able to implement it organizationally?
- Can you enforce it with security tools and/or sanctions?
- Does it clearly define areas of responsibility for the users, administrators, and management?
- Is it flexible and adaptable to changing environments?

RFC 2196 - http://www.ietf.org/rfc/rfc2196.txt

What Are You Protecting?

- Identify Critical Assets
  - Hardware, software, data, people, documentation
- Place a Value on the Asset
  - Intangible asset – importance or criticality
  - Tangible asset – replacement value, training costs and/or immediate impact of the loss
- Determine Likelihood of Security Breaches
  - What are threats and vulnerabilities ?

Impact and Consequences

- Data compromise
  - Stolen data
  - can be catastrophic for a financial institution
- Loss of data integrity
  - Negative press or loss or reputation (bank, public trust)
- Unavailability of resources
  - The average amount of downtime following a DDoS attack is 54 minutes
  - The average cost of one minute of downtime due to DDoS attack is $22,000*

* Based on a Ponemon Institute study (2012)

Risk Mitigation vs Cost

Risk mitigation: **the process of selecting appropriate controls to reduce risk to an acceptable level.**

The *level of acceptable risk* is determined by comparing the risk of security hole exposure to the cost of implementing and enforcing the security policy.

Assess the cost of certain losses and do not spend more to protect something than it is actually worth.

Will I Go Bankrupt ?

Is it an embarrassment ?

Evolution of Attack Landscape

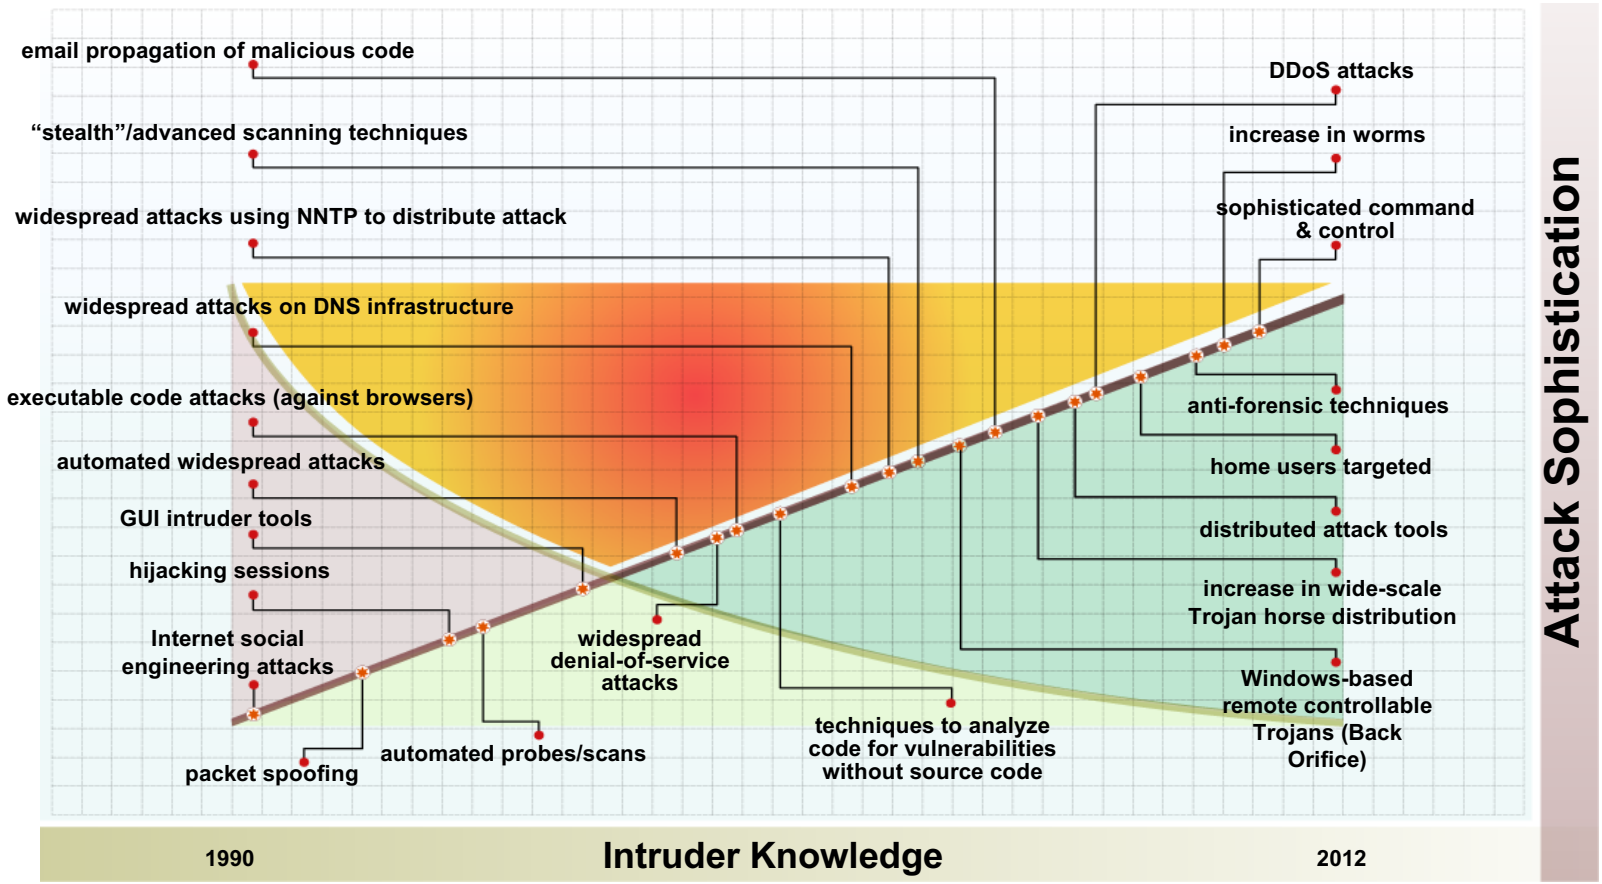

Attack Motivation

- Criminal
  - Criminal who use critical infrastructure as a tools to commit crime
  - Their motivation is money
- War Fighting/Espionage/Terrorist
  - What most people think of when talking about threats to critical infrastructure
- Patriotic/Principle
  - Large groups of people motivated by cause - be it national pride or a passion aka Anonymous

Attack Motivation

- Nation States want SECRETS
- Organized criminals want MONEY
- Protesters or activists want ATTENTION
- Hackers and researchers want KNOWLEDGE

Source: NANOG60 keynote presentation by Jeff Moss, Feb 2014

The Threat Matrix

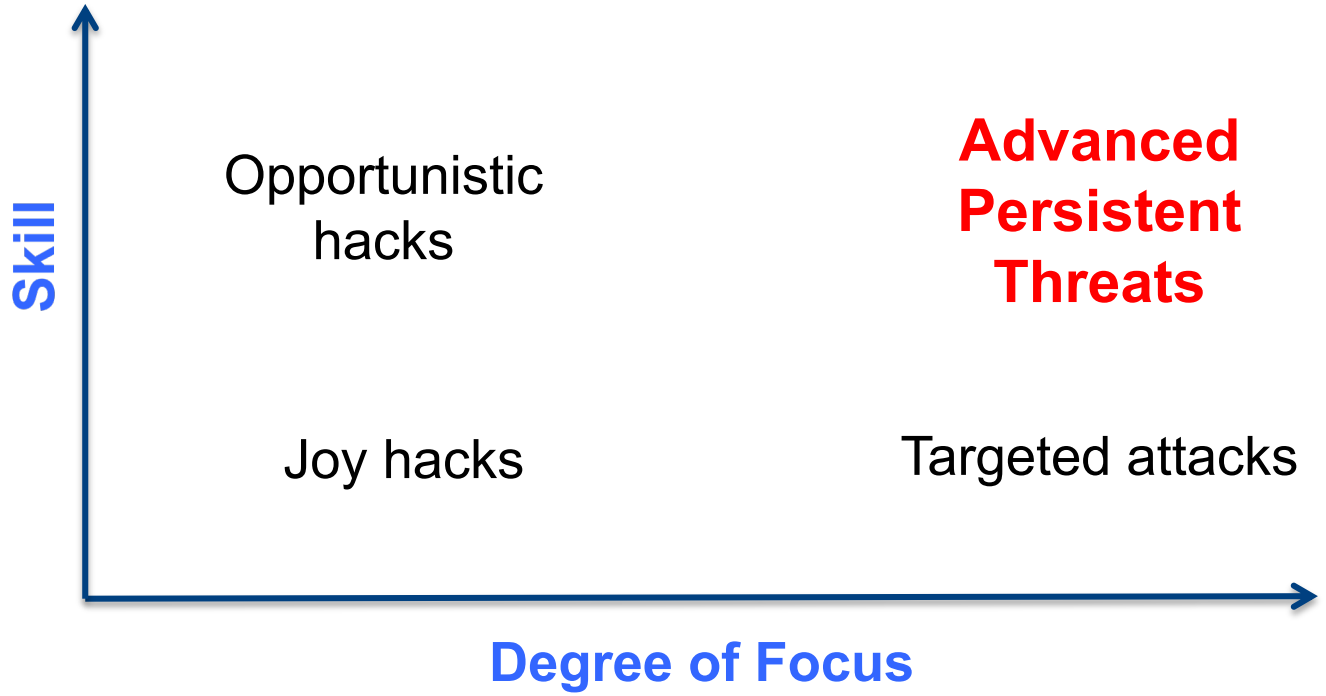

Attack Sources

- Active attack involves writing data to the network. It is common to disguise one's address and conceal the identity of the traffic sender.
- Passive attack involves only reading data on the network. Its purpose is breach of confidentiality.

Active Attacks	Passive Attacks
Denial of Service attacks Spoofing Man in the Middle ARP poisoning Smurf attacks Buffer overflow SQL Injection	Reconnaissance Eavesdropping Port scanning

Source: RFC 4778

Attack Sources

- On-path vs. Off-path
  - On-path hosts can read, modify, or remove any datagram transmitted along the path
  - Off-path hosts can transmit datagrams that appear to come from any hosts but cannot necessarily receive datagrams intended for other hosts
- Insider vs. outsider
  - What is definition of perimeter/border?
- Deliberate vs. unintentional event
  - Configuration errors and software bugs are as harmful as a deliberate malicious network attack

Source: RFC 4778

General Threats

- Masquerade
  - An entity claims to be another entity
- Eavesdropping
  - An entity reads information it is not intended to read
- Authorisation violation
  - An entity uses a service or resource it is not intended to use
- Loss or modification of information
  - Data is being altered or destroyed
- Denial of communication acts (repudiation)
  - An entity falsely denies its participation in a communication act
- Forgery of information
  - An entity creates new information in the name of another entity
- Sabotage
  - Any action that aims to reduce the availability and/or correct functioning of services or systems

Reconnaissance Attack

- Unauthorised users to gather information about the network or system before launching other more serious types of attacks
- Also called eavesdropping
- Information gained from this attack is used in subsequent attacks (DoS or DDoS type)
- Examples of relevant information:
  - Names, email address
    - Common practice to use a person's first initial and last name for accounts
  - Practically anything

Man-in-the-Middle Attack

- Active eavesdropping
- Attacker makes independent connections with victims and relays messages between them, making them believe that they are talking directly to each other over a private connection, when in fact the entire conversation is controlled by the attacker
- Usually a result of lack of end-to-end authentication
- Masquerading – an entity claims to be another entity

Session Hijacking

- Exploitation of a valid computer session, to gain unauthorized access to information or services in a computer system.
- Theft of a "magic cookie" used to authenticate a user to a remote server (for web developers)
- Four methods:
  - Session fixation – attacker sets a user's session id to one known to him, for example by sending the user an email with a link that contains a particular session id.
  - Session sidejacking – attacker uses packet sniffing to read network traffic between two parties to steal the session cookie.

Denial of Service (DoS) Attack

- Attempt to make a machine or network resource unavailable to its intended users.
- Purpose is to temporarily or indefinitely interrupt or suspend services of a host connected to the Internet
- Saturating the target with external communications requests (server overload)
  - May include malware to max out target resources, trigger errors, or crash the operating system
- DDoS attacks are more dynamic and comes from a broader range of attackers
- Can be used as a redirection and reconnaissance technique

Reflected Denial of Service (rDoS)

- Involves sending forged requests to hundreds of machines with replies directed to a victim server
- Attacker modifies the source IP address (spoofing)
- Replies are expected to be much bigger than the request
- DNS is used for this due to its lack of source validation

Summary - Most Common Threats and Attacks

- Unauthorized access – insecure hosts, cracking
- Eavesdropping a transmission – access to the medium
  - Looking for passwords, credit card numbers, or business secrets
- Hijacking, or taking over a communication
  - Inspect and modify any data being transmitted
- IP spoofing, or faking network addresses
  - Impersonate to fool access control mechanisms
  - Redirect connections to a fake server
- DOS attacks
  - Interruption of service due to system destruction or using up all available system resources for the service
  - CPU, memory, bandwidth

Attack Trends

- Key findings:
  - Largest DDoS attack at 600Gbps (down from 800Gbps in 2016)
  - Multiple attacks over 300Gbps
  - Hacktivism is top commonly perceived motivation behind attacks
  - Customers are the most common target of attacks (75%), with service infrastructure coming second (15%)

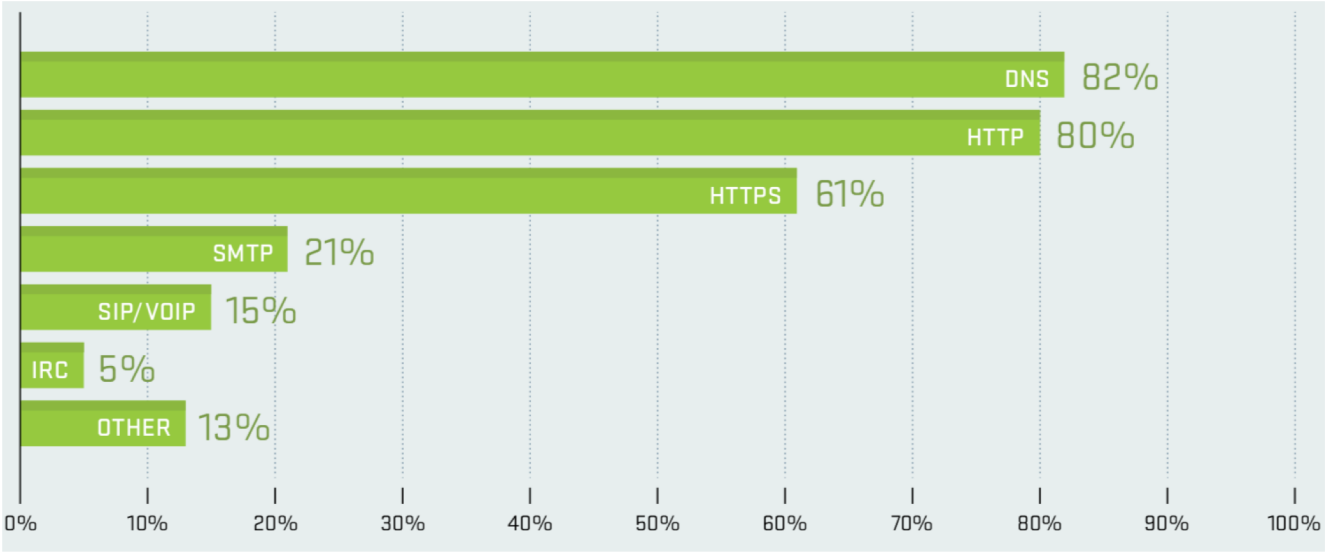

Source: NETSCOUT Arbor 13th Worldwide Infrastructure Security Report

Attack Trends

- Volumetric attacks were the most common, at 76% of all DDoS
  - DNS and NTP the most frequently used protocols
  - CharGen, SSDP and SNMP close behind

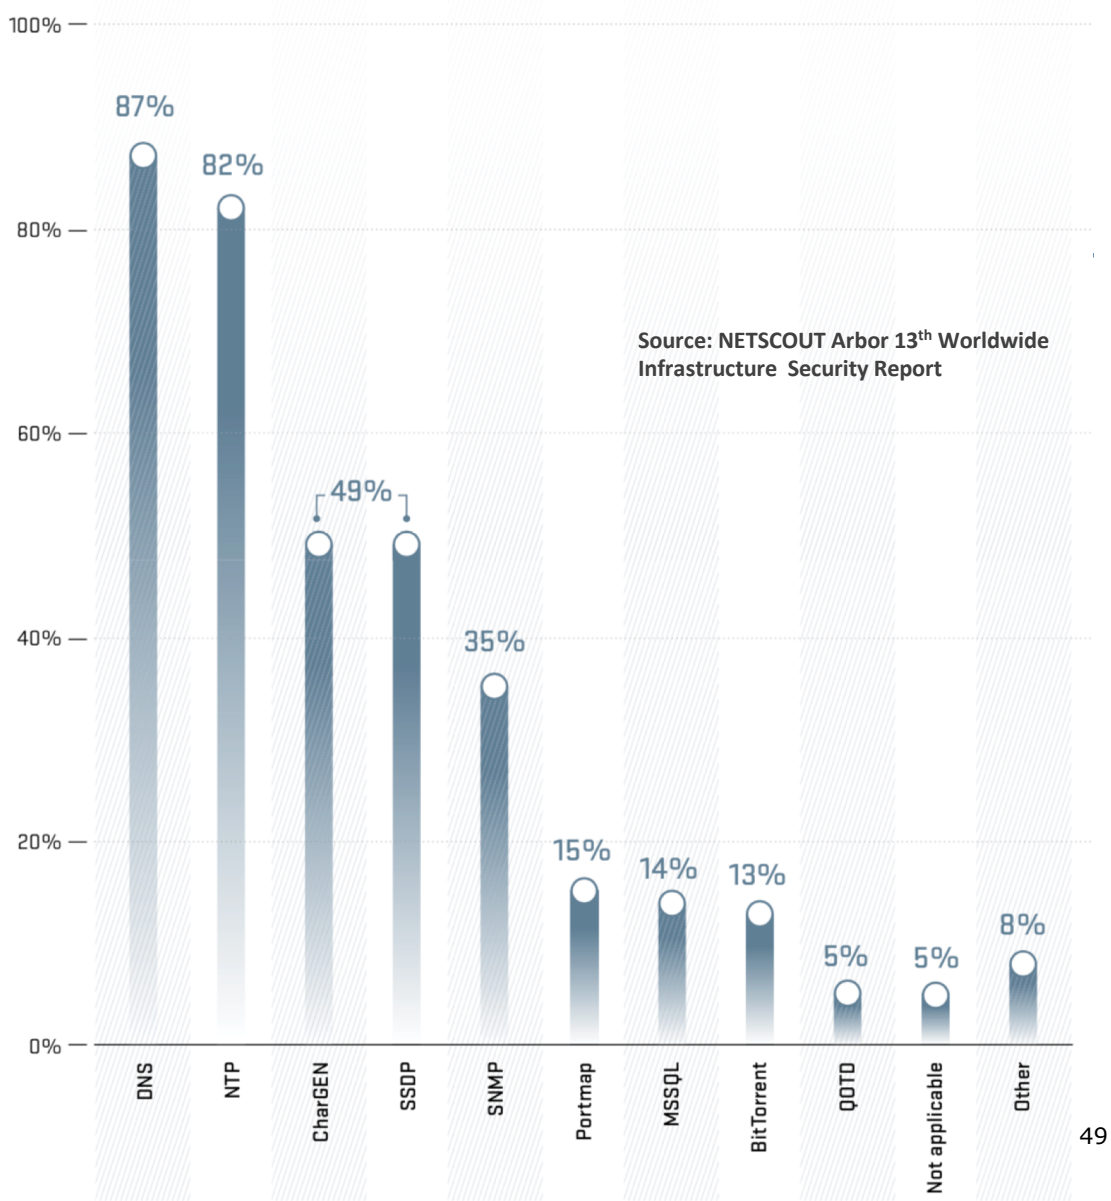

Attack Trends

- Downward trend in the use of application-layer attacks (only 12% in 2017)

Source: NETSCOUT Arbor 13th Worldwide Infrastructure Security Report

- "To launch significant DDoS layer 7 attack campaigns, attackers need to possess sophisticated skills. Few attackers are capable of these attacks, as it requires compromising servers and applications by the exploitation of vulnerabilities, and often requires code customization"

Source: Prolexic Q2 2014 Global DDOS Attack Report

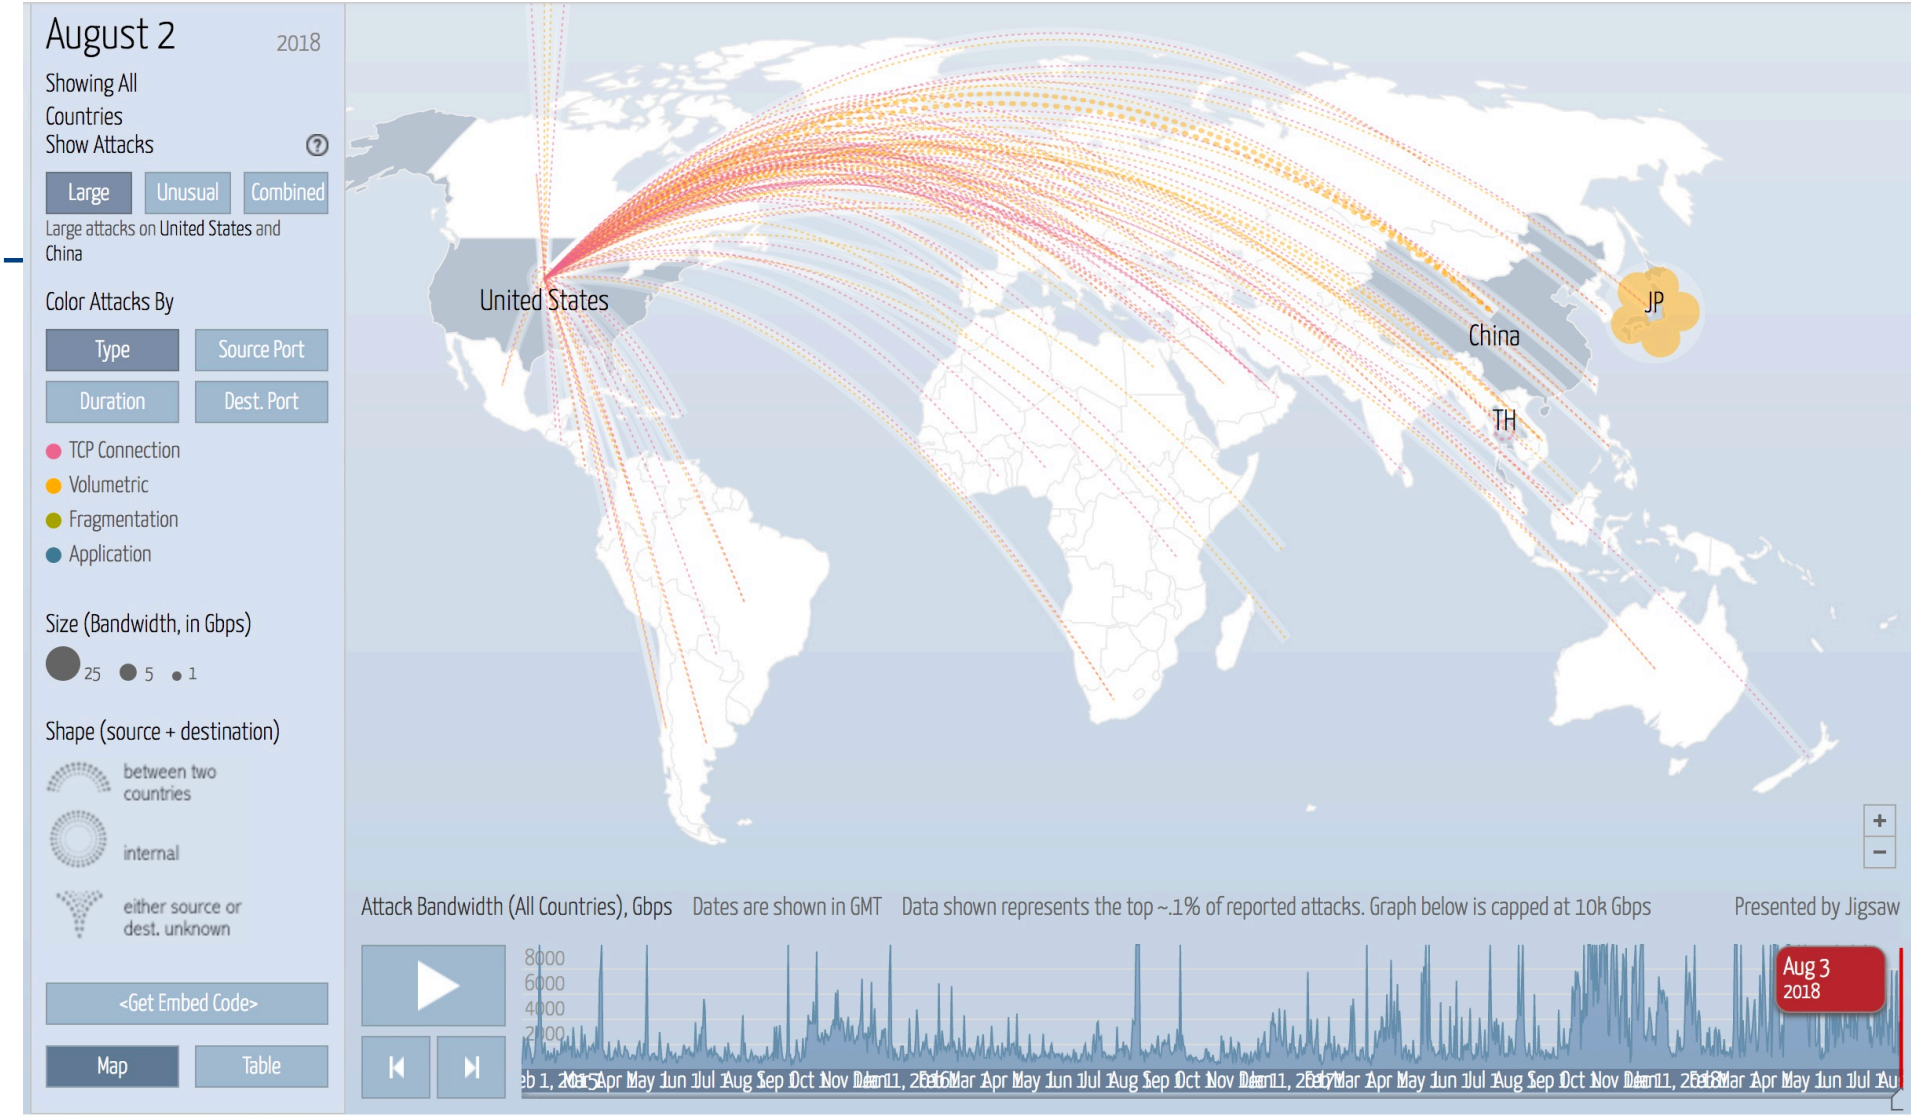

Global Map of DDoS Attacks

Source: http://www.digitalattackmap.com 4th August 2018

Web Attack Visualisation

Source: https://globe.akamai.com/

4th August 2018

Internet Security Introduction

ISP Workshops