Attack Prevention and Detection

David M. Nicol

Department of Energy Short Course on
Cyber Security for the Power Grid
Washington, DC

Monday, March 25, 2013

Overview

- Authentication
- Cryptography
- Intrusion Detection
- Firewalls
- NERC SIP

Authentication

Basics

- Authentication: binding of identity to subject
  - Identity is that of external entity (my identity, the Illini Union Bookstore, *etc*.)
  - Subject is computer entity (process, network connection, *etc*.)
- Two steps
  - Identification step: present identifier to security system. Registration
  - Verification step: Present or generate authentication information that corroborates the binding between entity and identifier

Establishing Identity

- One or more of the following
  - What entity knows (*e.g.* password, private key)
  - What entity has (*e.g.* badge, smart card)
  - What entity is (*e.g.* fingerprints, retinal characteristics)
  - What entity does (e.g., voice pattern, handwriting, typing rhythm)
  - Where entity is (*e.g*. In front of a particular terminal)
- Example: Credit card transaction
- Multi-factor authentication
  - Use multiple elements to prove identity

Password-based Authentication

- External entity is bound to system ID (user account)
- Authentication Step
  - External entity presents password
  - System compares with previously stored password
  - If password matches, system starts process with bound ID
- Later access control decisions made against ID
- Privilege decisions made against ID

Password Vulnerabilities

- Password systems widely used, but very vulnerable
  - Offline dictionary attack
  - Specific account attack
  - Workstation hijacking
  - Sticky notes
  - Password reuse
  - Social engineering
  - Electronic monitoring

Password Storage

- Store as cleartext
  - If password file compromised, *all* passwords revealed
- Encipher file
  - Need to have decipherment, encipherment keys in memory
  - Reduces to previous problem
- Store one-way hash of password
  - If file read, attacker must still guess passwords or invert the hash

Dictionary Attacks

- Trial-and-error from a list of potential passwords
  - *Off-line (type 1)*: know functions and registered information, and repeatedly try different guesses $g \in A$ until the list is done or passwords guessed
    - Examples: *crack*, *john-the-ripper*
  - *On-line (type 2)*: have access to verification functions. Try guesses until one succeeds.
    - Examples: trying to log in by guessing a password

Preventing Attacks

- How to prevent this:
  - Hide information so that either authentication input, authentication functions, or stored verification information cannot be found. Prevents obvious attack from above
    - Example: UNIX/Linux shadow password files
      - Hides *c*'s
  - Block access to all verification methods
    - Prevents attacker from knowing if guess succeeded
    - Example: preventing *any* logins to an account from a network
      - Prevents knowing results of verification function or accessing verification function.

Approaches: Password Selection

- Random selection
  - Any password from *A* equally likely to be selected
  - See previous example
  - Make sure it's random!
- Pronounceable passwords
- User selection of passwords

User Password Education

- Use the first letter of each word in a phrase
  - "My dog's first name is Rex." becomes "MdfniR"

Token Based Authentication

- Memory Cards
  - Stores data like an ATM card
  - Does no computation
  - Generally combined with PIN (2-factor)
- Smart Card
  - Storage and computation

Token-based Authentication Protocols

- Static:
  - User authenticates to memory card, and memory card authenticates to system
- Dynamic Password Generator:
  - System periodically creates and displays new password
  - User enters current password
- Challenge response:
  - See next slide

Challenge-Response

- User and system share a secret function
- User proves knowledge of secret function by answering challenge

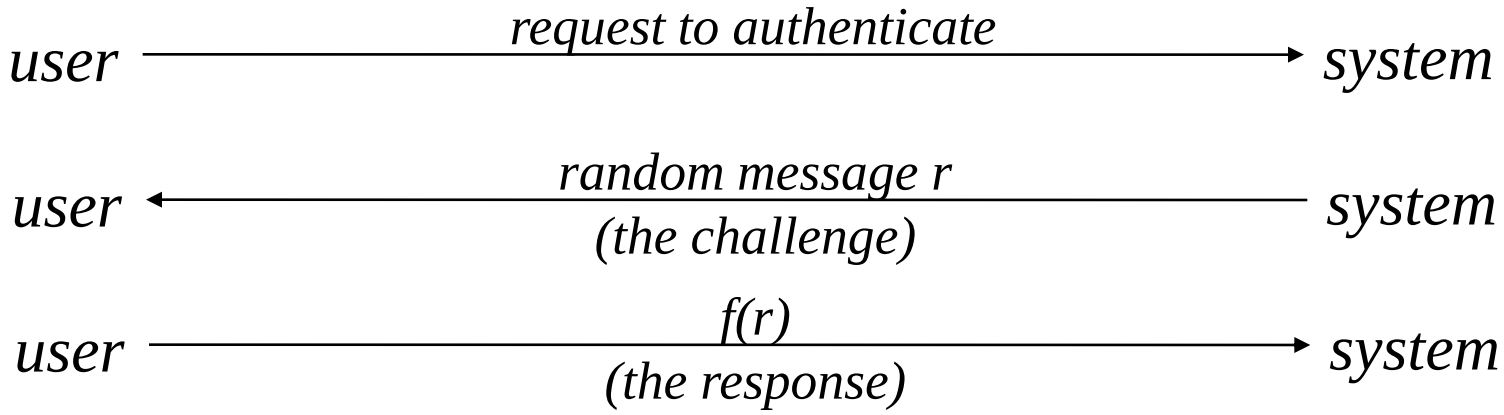

Token-based Authentication

- Something you have
- Memory Cards
  - No computation on the card
  - Need special reader to pull data off the card
  - Need pin to decrypt data off of card
  - E.g., ATM card or debit card
- By adding PIN (something you know) you get multi-factor authentication

Token Based Authentication

- Smart Card
  - Computation on the card
  - Plug in with USB or wireless communication (credit card)
- Authentication options
  - Static – equivalent to memory card
  - Dynamic password generator – generates a unique password every minute.
  - Challenge response

Biometrics

- Automated measurement of biological, behavioural features that identify a person
  - Fingerprints: optical or electrical techniques
    - Maps fingerprint into a graph, then compares with database
    - Measurements imprecise, so approximate matching algorithms used
  - Voices: speaker verification or recognition
    - Verification: uses statistical techniques to test hypothesis that speaker is who is claimed (speaker dependent)
    - Recognition: checks content of answers (speaker independent)

Other Characteristics

- Can use several other characteristics
  - Eyes: patterns in irises unique
    - Measure patterns, determine if differences are random; or correlate images using statistical tests
  - Faces: image, or specific characteristics like distance from nose to chin
    - Lighting, view of face, other noise can hinder this
  - Keystroke dynamics: believed to be unique
    - Keystroke intervals, pressure, duration of stroke, where key is struck
    - Statistical tests used

Cryptography

Classical Cryptography

- Sender, receiver share common key
  - *symmetric cryptography*
    - Keys may be the same, or trivial to derive from one another
- Two basic types
  - Transposition ciphers
  - Substitution ciphers
  - Combinations are called *product ciphers*

Transposition Cipher

- Rearrange letters in plaintext to produce ciphertext
- Example : Reverse every group of four non-space characters

THE WINTER OF OUR DISCONTENT

Becomes

WEHTETNIOFORIDRUNOCNTNET

- More generally known as ‘permutation’

Substitution Ciphers

- Change characters in plaintext to produce ciphertext
- Example (Cæsar cipher)
  - Plaintext is HELLO WORLD
  - Change each letter to the third letter following it (X goes to A, Y to B, Z to C)
    - Key is 3, usually written as letter ‘D’
  - Ciphertext is KHOOR ZRUOG

Substitution Ciphers

- Modern crypto systems substitute blocks of characters, in conjunction with other operations
- ‘S-box’ (substitution box)

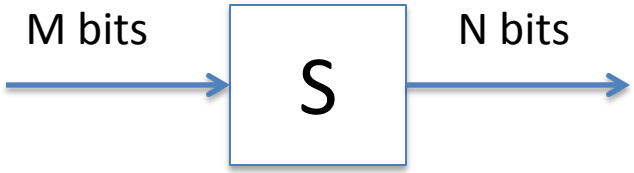

- Modern crypto systems substitute blocks of characters, in conjunction with other operations
- ‘S-box’ (substitution box)

Block Cipher

View a message as a sequence of data blocks of the same size

$$B_0 B_1 B_2 B_3 B_4 B_5 \ldots$$

Encryption of a block using key K:

$$E_K(B_i)$$

Sequence of keys

Encrypted s

$$K_0 \; K_1 \; K_2 \; K_3 \; K_4 \; K_5 \ldots$$

$$E_{K_0}(B_0) E_{K_1}(B_1) E_{K_2}(B_2) E_{K_3}(B_3) E_{K_4}(B_4) E_{K_5}(B_5) \ldots$$

One-time Pad

Keys are random, jointly independent

- Means knowledge of any subset gives no information about any particular key
- “One-time pad” means that keys are not re-used
- Impossible to break except by brute force

Hard to realize, approximations

- “Book cipher”
- Enigma machine

AES (Advanced Encryption Standard)

US NIST issued call for ciphers in 1997

Requirements

- Private key symmetric block cipher
  - 128-bit data, 128/192/256-bit keys
- Stronger & faster than Triple-DES
- Active life of 20-30 years (+ archival use)
- Provide full specification & design details
- Both C & Java implementations
- NIST have released all submissions & unclassified analyses

AES Evaluation Criteria

- general security
- software & hardware implementation ease
- implementation attacks
- flexibility (in en/decrypt, keying, other factors)

AES Overview

Everything organized around a ‘block matrix’ of bytes as the basic unit of encryption (128 bits, below)

$a_{0,0}$	$a_{0,1}$	$a_{0,2}$	$a_{0,3}$
$a_{1,0}$	$a_{1,1}$	$a_{1,2}$	$a_{1,3}$
$a_{2,0}$	$a_{2,1}$	$a_{2,2}$	$a_{2,3}$
$a_{3,0}$	$a_{3,1}$	$a_{3,2}$	$a_{3,3}$

AES Overview

Steps involve

a) byte-by-byte substitution

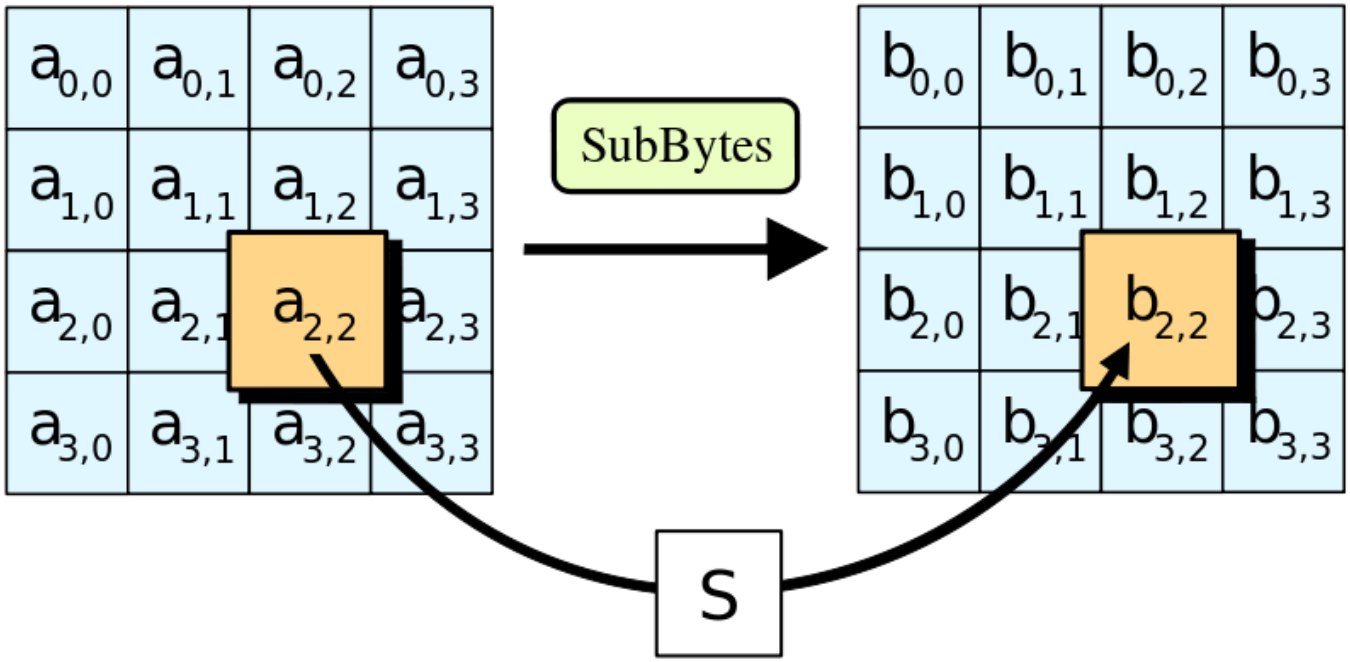

AES Overview

Steps involve

b) Shifting rows

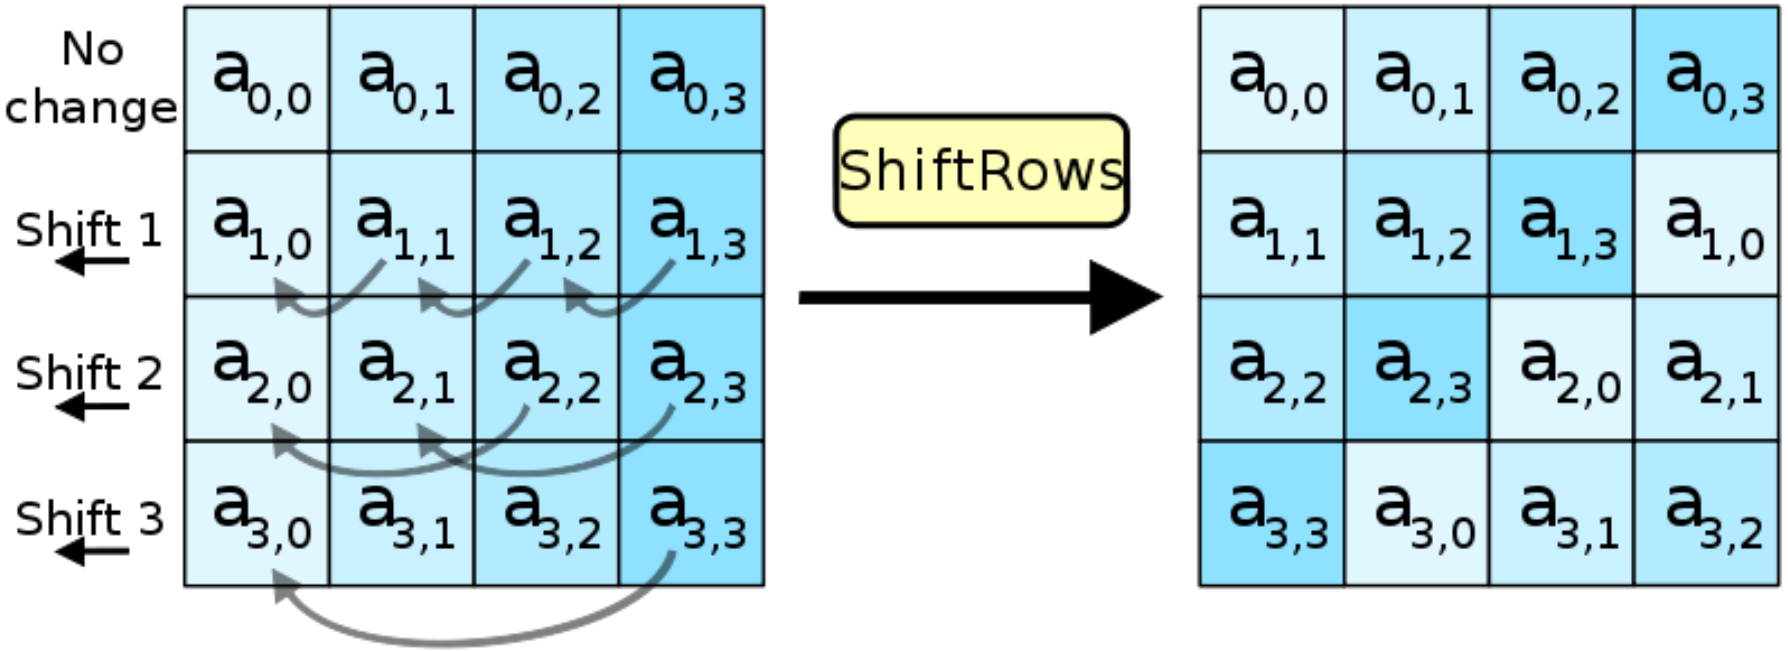

AES Overview

Steps involve

c) Mixing columns (fancy math: multiplication over Galois field)

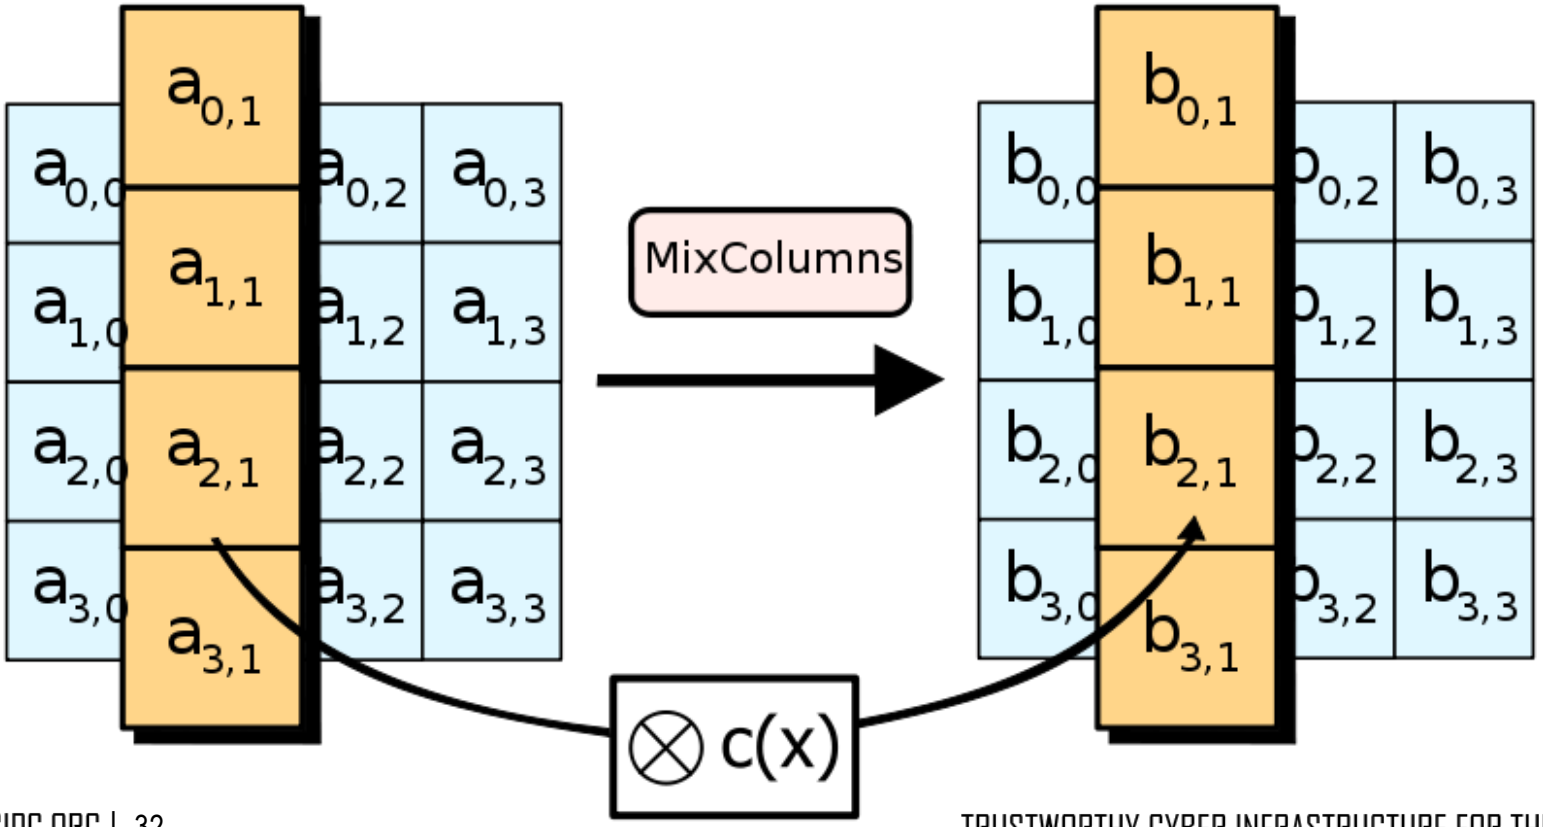

AES Overview

Steps involve

d) Use of key (different key per ‘round’

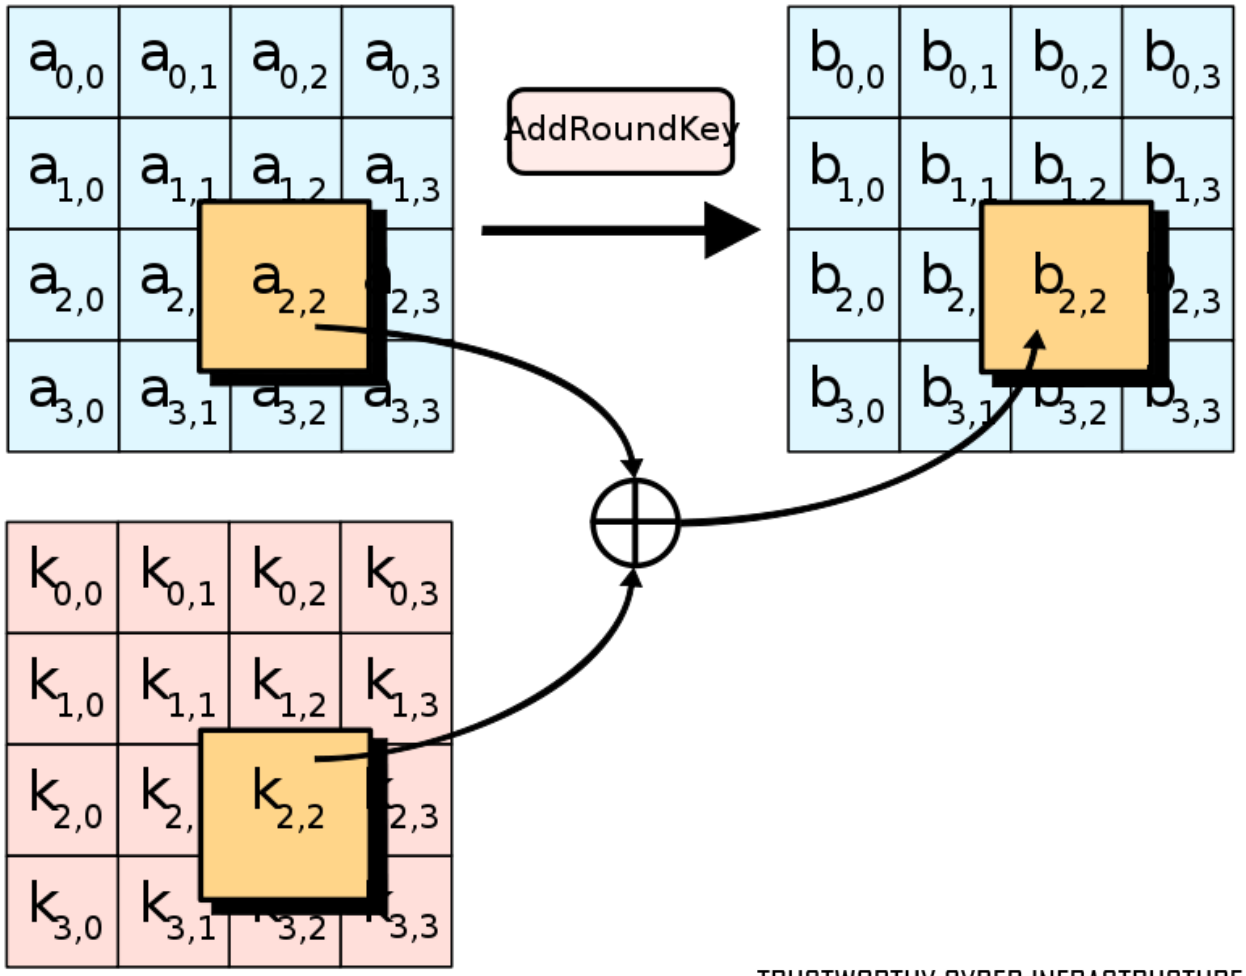

AES Overview

This sequence of steps applied in each of several ‘rounds’

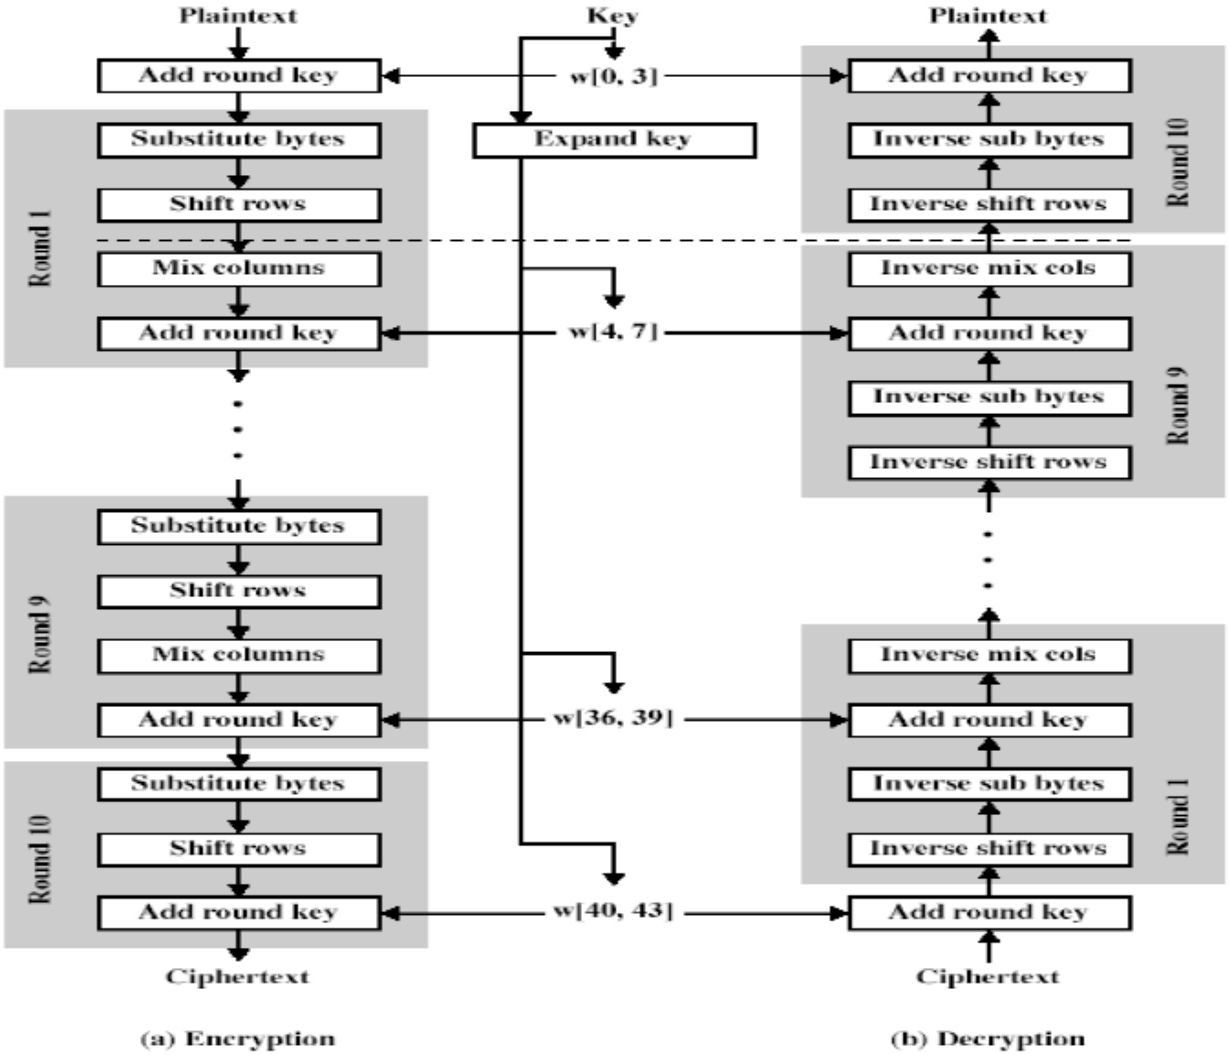

Public Key Cryptography

Key K = (e,d) has two parts

- ‘e’, the “public” part
- ‘d’, the private part

Alice wants to send message m to Bob

- Bob’s key is (eBob,dBob)

- She looks up on a database Bob’s public key, eBob
- She encrypts (particular function E) m using eBob, E(eBoB,m)
  - Can be transmitted in the clear
- Bob receives E(eBob,m), applies decode function D using private key dBob:
  - D(dBob,E(eBob,m)) = m

Alice and Bob share NO secret information!

Public Key Cryptography

Can be used to prove possession of a private key:

- A message is 'signed' with Alice's private key dAlice
  - D(dAlice,m)
  - Note use of D, not E, but D==E in some systems
- Bob (and anyone else) receives message (m, D(dAlice,m)) and can look up Alice's public key eAlice
  - Compute E(eAlice,D(dAlice,m)) = m, and compare with m in message
- What this proves:
  - (a) that the signature computed was computed on m
    - Providing a proof of integrity
  - (b) that the signer had possession of Alice's private key
    - Is this proof that the signer was Alice?
- RSA most common algorithm of this type

Public Key Cryptography

- In practice public key crypto is computationally much more expensive than symmetric
- It is used principally to
  - Encode and exchange symmetric 'session keys', after which the crypto is done using AES or some other symmetric scheme
  - Sign a digest of a message, as a means of integrity
- A 'one-way' hash function h takes an arbitrarily long message m, and computes h(m) whose result size is fixed, e.g., 180 or 256 bytes
  - 'one-way' means that knowing h(m) it is computationally difficult to discover m
- Given (m,D(private,h(m))) the receiver can
  - Compute h(m)
  - Compare computed h(m) with E(public,D(private,h(m)))

Public Key Infrastructure

- PKI is a framework for using and managing public key crypto systems
- Most prevalent use is with digital certificates
- Problem : Bob receives signed message, reported from Alice. Why should he believe it is from Alice?
- A (digital) certificate is issued by a certificate authority who
  - Verifies that a requester claiming to be Alice, is indeed Alice
  - Creates a digital document that
    - States this certificate asserts that Alice's identity has been verified
    - Contains a public key to be used in communication with Alice
    - Contains a signed hash of the certificate, signed by the private key of the certificate authority
- So if Bob trusts the certificate authority (to what?), he will trust the public key that the certificate holds

Use of certificates

- SSL connection
  - Client wishes to communicate sensitive information to a server (e.g., credit card)
  - Server offers certificate with its public key
    - If client trusts the certificate and the certificate issuer, sensitive information can be encrypted with server's public key
- Integrity of software
  - e.g., driver for some device
  - Certificate issued for software, contains h(C), signed by private key of issuer
    - System looks up public key of issuer, recomputes h(C), decodes version contained in cert, accepts if comparison checks out

Intrusion Detection

Intrusion Detection System

- Device or software that monitors network or activities
  - For things know or suspected of being malicious
  - Policy violations
- NIDS (Network Intrusion Detection System)
  - Monitor traffic via connection to hub or switch, placed for best visibility
    - Example: SNORT
- HIDS (Host-based Intrusion Detection System)
  - Watches system calls, log files, file-system modifications
    - Example: Tripwire

Intrusion Detection System

Passive System

– Raises alerts, but takes no actions

Reactive System

– When “enough” evidence present of intrusion, some action taken, e.g.
  - Reset connection
  - Isolate LAN
  - Block traffic from suspected source
– Any reactive system has potential for being vector for attack designed to repeatedly trigger
  - Because legitimate use is impacted by reaction

Statistical Anomaly Detection

Idea is to observe statistics of behavior, and react

- Assumes some notion of "normal" is known

Difficulties

- "normal" can have high variance, making it easier for abnormal to look normal
- Attributes of behavior have to be specified
  - The more the attributes the better the ability to differentiate, but the less data is there is train on
- What's the difference between "abnormal" and "rare"???

Signature Based IDS

System examines packets (particularly coming into host) and looks for malware

- Needs a database of “signatures”
- Most common defense against viruses, but
  - There’s a lag between when new malware appears, and when it can be identified and put in the DB
  - Malware can hide by auto-mutation to change signatures.

Firewalls

Firewall Goal

- Insert *after the fact security* by wrapping or interposing a filter on network traffic

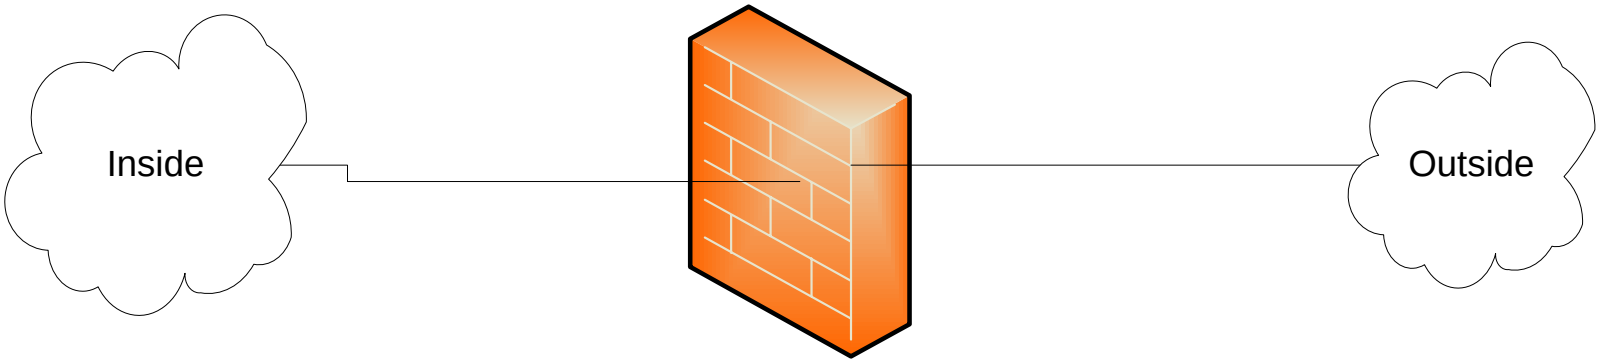

Firewall Requirements

- All traffic between network section A and network section B (and visa versa) must pass through the firewall (or a consistently controlled set of firewalls)
- Only authorized traffic (as specified by the security policy) is allowed to pass
- The firewall itself is immune to penetration

“Typical” corporate network

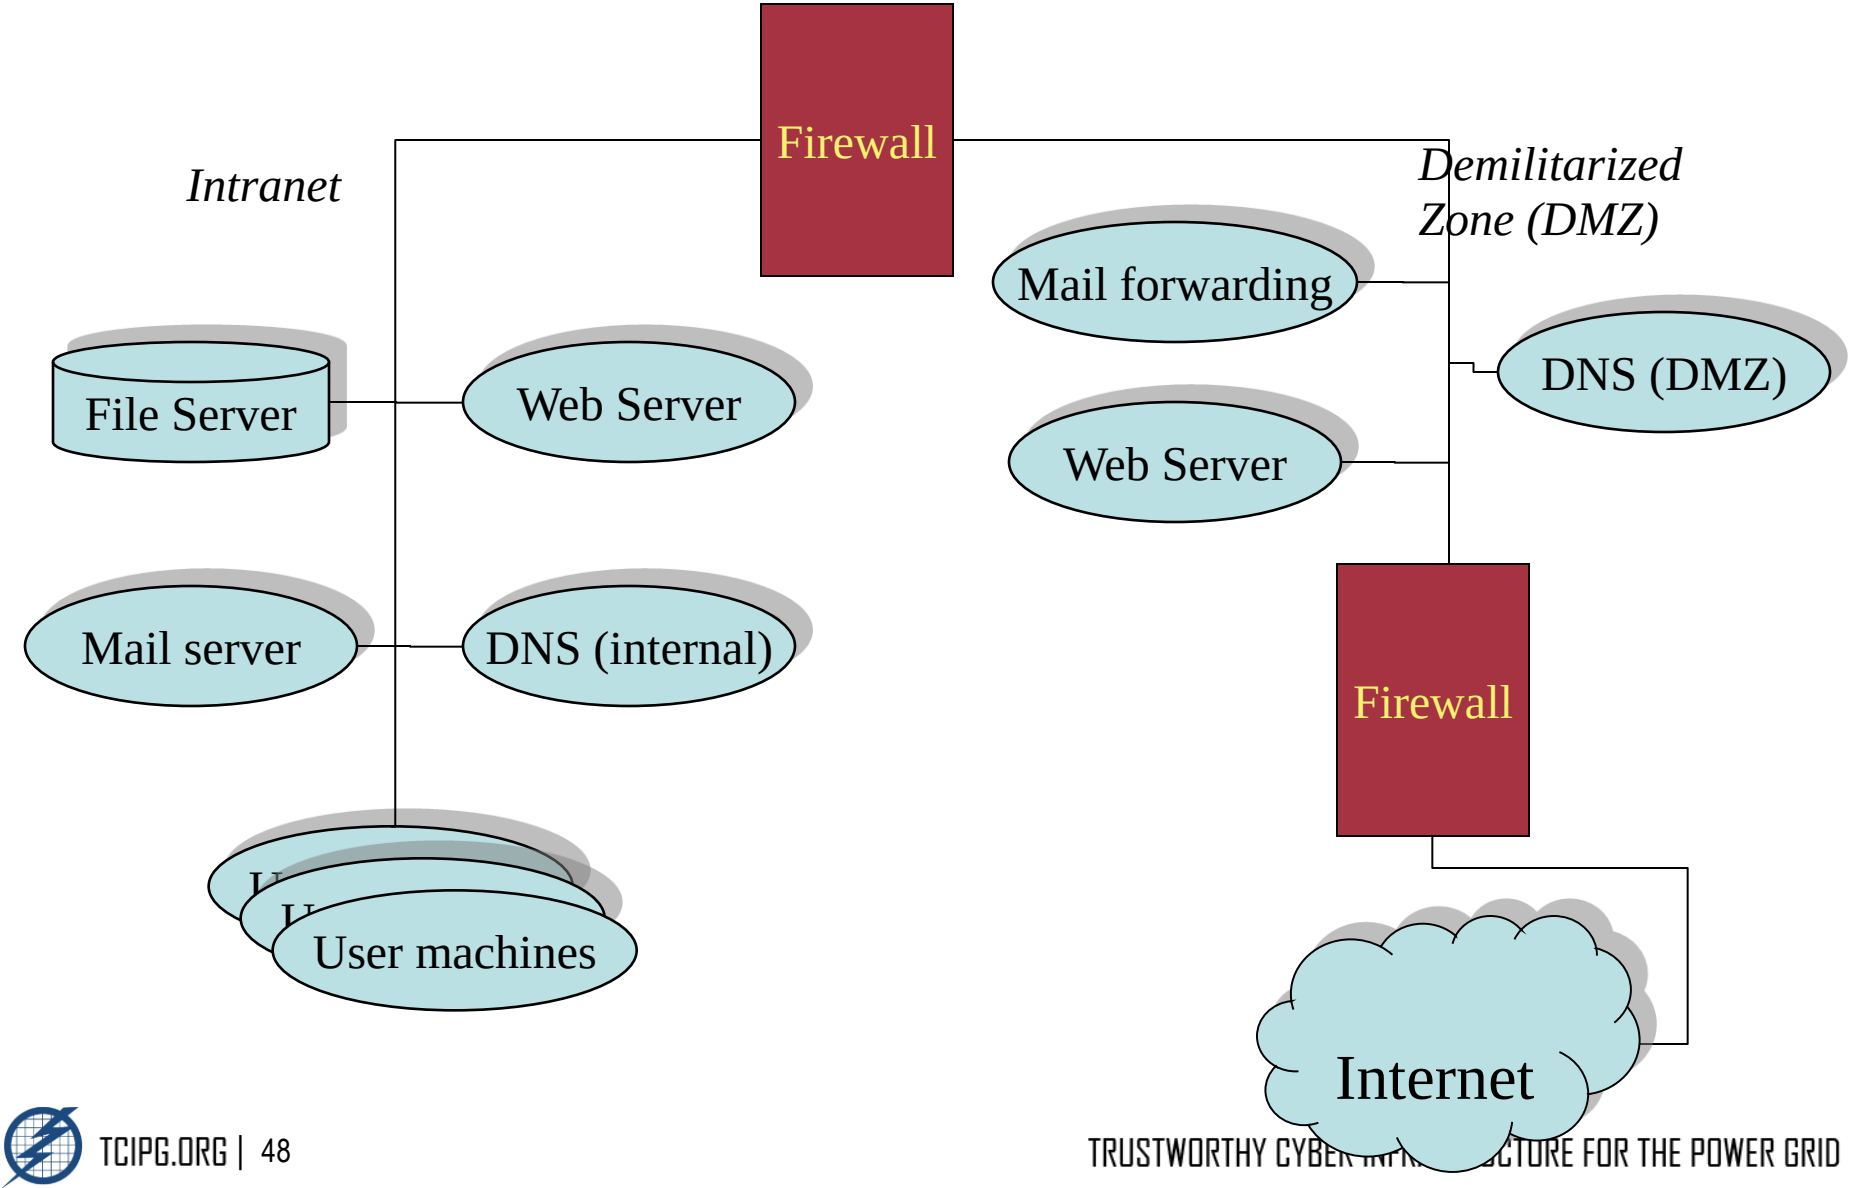

Packet Filter Firewall

- IP packet contains
  [src ip, src port, dst ip, dst port, protocol]
  - (ip,port) is standard Internet address
  - Protocol describes how the packet is handled by sender and receiver

Packet Filter

- Can block traffic based on source and destination address, ports, and protocol
- Does not look at data contents

Rule Scenario

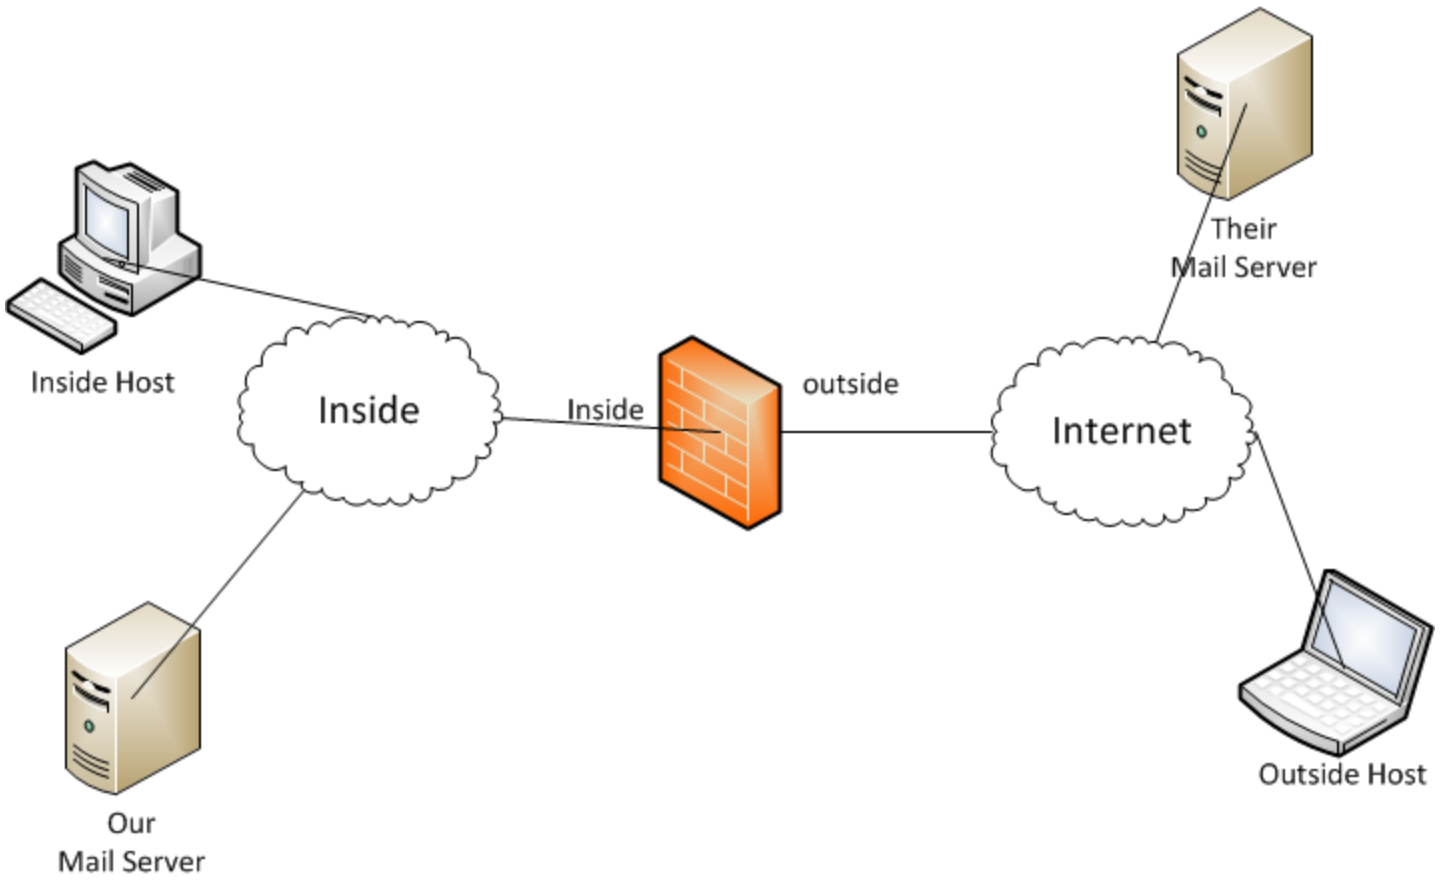

Example Packet Filter Rules

- Rules attached to outside interface

Action	Source Addr	Src port	Dest Addr	Dest Port	Protocol	Comment
Block	Outside host	*	*	*	*	Don't trust
Allow	*	*	Our Mail Server	25	TCP	Allow mail traffic

Action	Source Addr	Source Port	Dest Addr	Dest Port	Protocol	Comment
Block	*	*	Outside host	*	*	Don't trust
Allow	Our Mail Server	25	*	*	TCP	Allow Mail traffic

A Better Example

- Rules attached to inside interface

Action	Source Addr	Source Port	Dest Addr	Dest Port	Proto	Comment
Allow	Inside networks	*	*	25	TCP	Allow traffic to all mail servers

Action	Source Addr	Source Port	Dest Addr	Dest Port	Proto	Flags	Comment
Allow	*	25	Inside networks	*	TCP	ACK	Allow return traffic from all mail servers

Application Proxy Firewall

- Firewall software runs in application space on the firewall
- The traffic source must be aware of the proxy and add an additional header
  - Now transparent proxy support is available (TPROXY)
- Leverage basic network stack functionality to sanitize application level traffic
  - Block java or active X
  - Filter out “bad” URLs
  - Ensure well formed protocols or block suspect aspects of protocol

Ingress and Egress Filtering

- Ingress filtering
  - Filter out packets from invalid addresses before entering your network
- Egress filtering
  - Filter out packets from invalid addresses before leaving your network

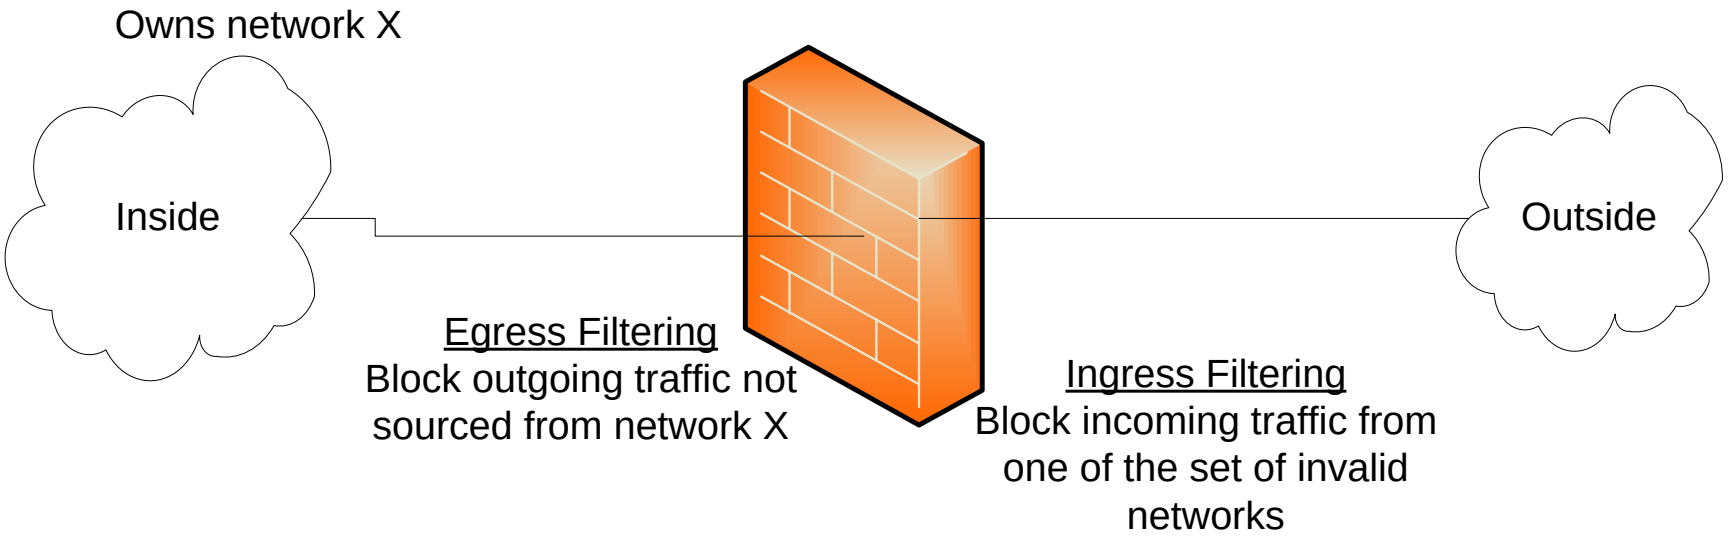

Limits to firewalls

- Cannot analyze encrypted traffic
  - Beyond header information
- Relies on port as indicator of service
  - Newer firewalls dynamically analyze traffic to determine protocol
- Tracking IP addresses instead of people
- Management is complex

NERC CIP Standards

NERC

FERC contracts with North American Electric Reliability Corporation for

- Development of reliability standards for bulk electric power systems

- Compliance enforcement of NERC standards
  - Monitoring, audits, investigations
  - Financial penalties for non-compliance

NERC Auditing

NERC conducts periodic, independent assessments of reliability and adequacy of utilities

- Development of reliability standards for bulk electric power systems

- Compliance enforcement of NERC standards
  - Monitoring, audits, investigations
  - Financial penalties for non-compliance

NERC Standards

- Developed with industry cooperation
- Reviewed and approved by NERC Board of Trustees
- Approved by FERC
  - Then become law

NERC Standards Overview

There are 14 different standards domains

Here we are interested in those pertaining to

- Critical Infrastructure Protection
  - Identify control elements critical to operation, and protection of access to them
  - Electronic Security Barrier (CIP—05)
    - Networking boundary behind which all critical cyber assets reside
    - Assess through the ESP

CIP 005

- CIP 005 presents challenges in both compliance and audit

- Requirement R1 calls for the documentation of an Electronic Security Perimeter and the access points to the ESP
  - any connection that crosses the ESP from outside (e.g. dial-up modem) to device inside

  - Endpoints of connections between different ESPs

- Requirement R1 requires documentation of all connections within ESP

Typical Utility Installation

Process Control networks are connected in enterprise systems

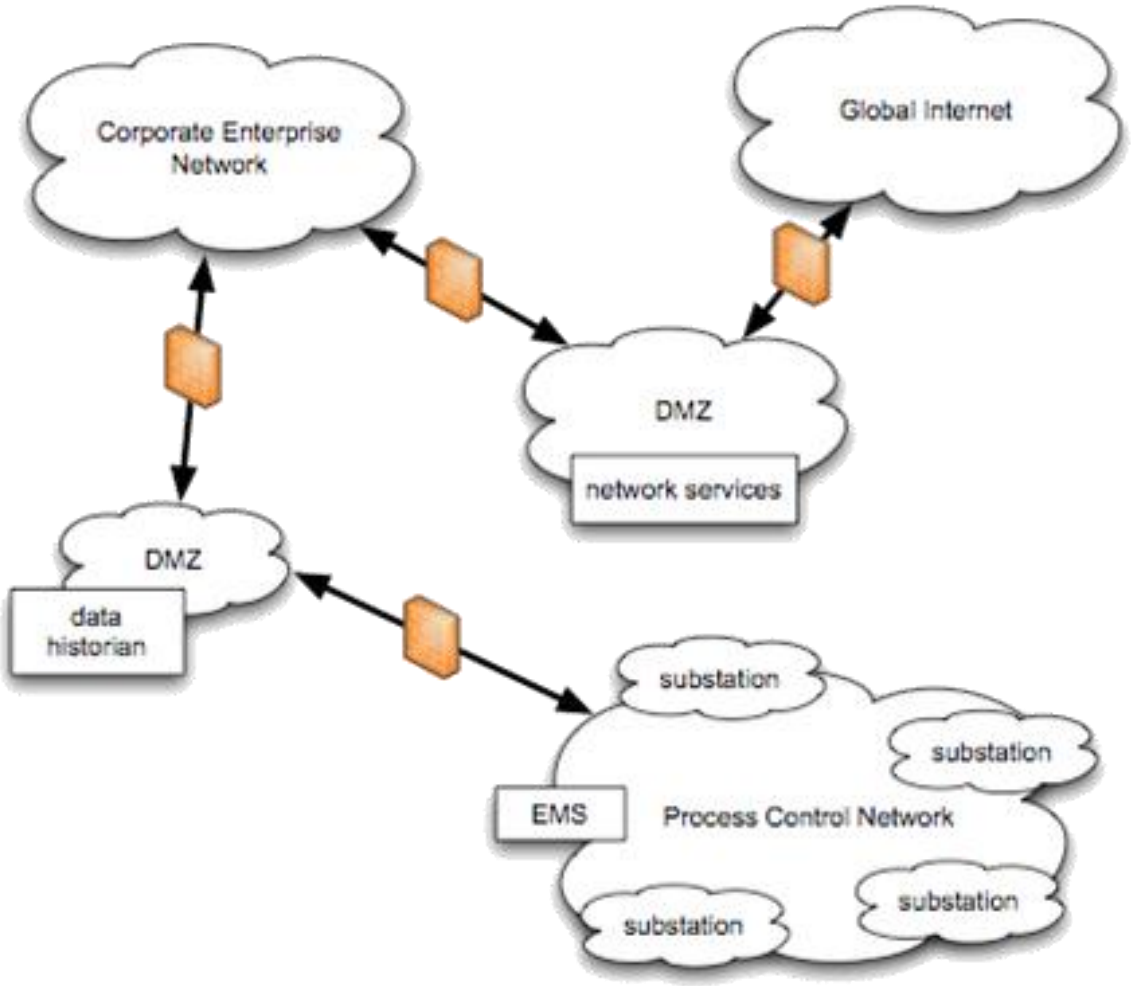

ESP

- Many utilities use a combination of diagrams and the configurations of their firewalls to define the ESP and security controls surrounding the ESP

- The utility needs to document organizational processes and technical and procedural mechanisms for control of electronic access

Vulnerability Assessment

- The vulnerability assessment requirement of CIP 005 calls for a review of controls

- The review of connections available in a firewall are not easily found without the use of a tool or the expenditure of significant man hours.

The need for a tool

- Utilities primarily use firewalls to establish access points to their critical systems
- Reviewing long and often complex firewall configurations (sometimes exceeding thousands of lines) is time consuming
- There is no automated method to verify the diagram of the ESP
- Audits are typically scheduled to last only a matter of days and auditors need a method to process this complex data without connecting to the network

Our tool : Network Access Tool (NetAPT)

- From only the configuration files of firewalls, NetAPT
  - Finds the devices and how the network is connected
  - Displays network graphically
  - Finds all connections that the firewalls allow

NetAPT

- NetAPT user can
  - Find all ways that a given device can be reached
  - Find all ways that a given device can reach others
  - Find all connections that use a certain protocol

Example of an EMS network

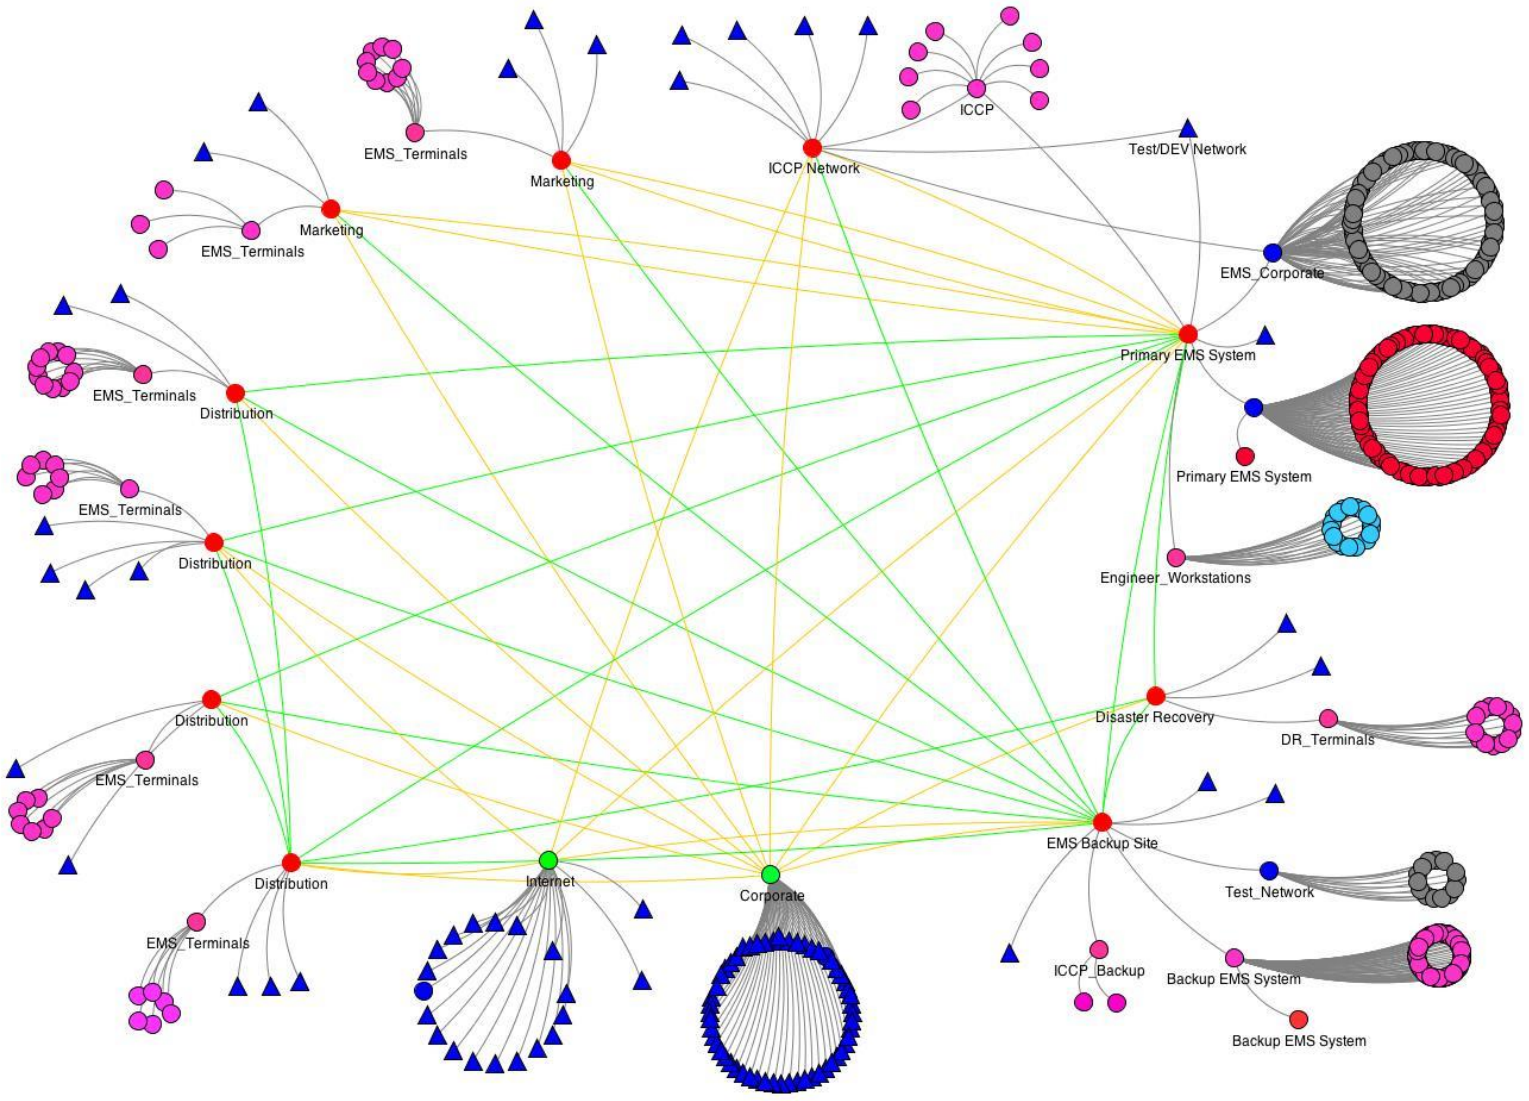

Analysis

- Analysis can be performed using the global policies established in the tool
- Analysis will show all incoming traffic allowed through the access point into the critical network
- Analysis can verify that all Critical Cyber Assets are accounted for and there are no extraneous CCA's

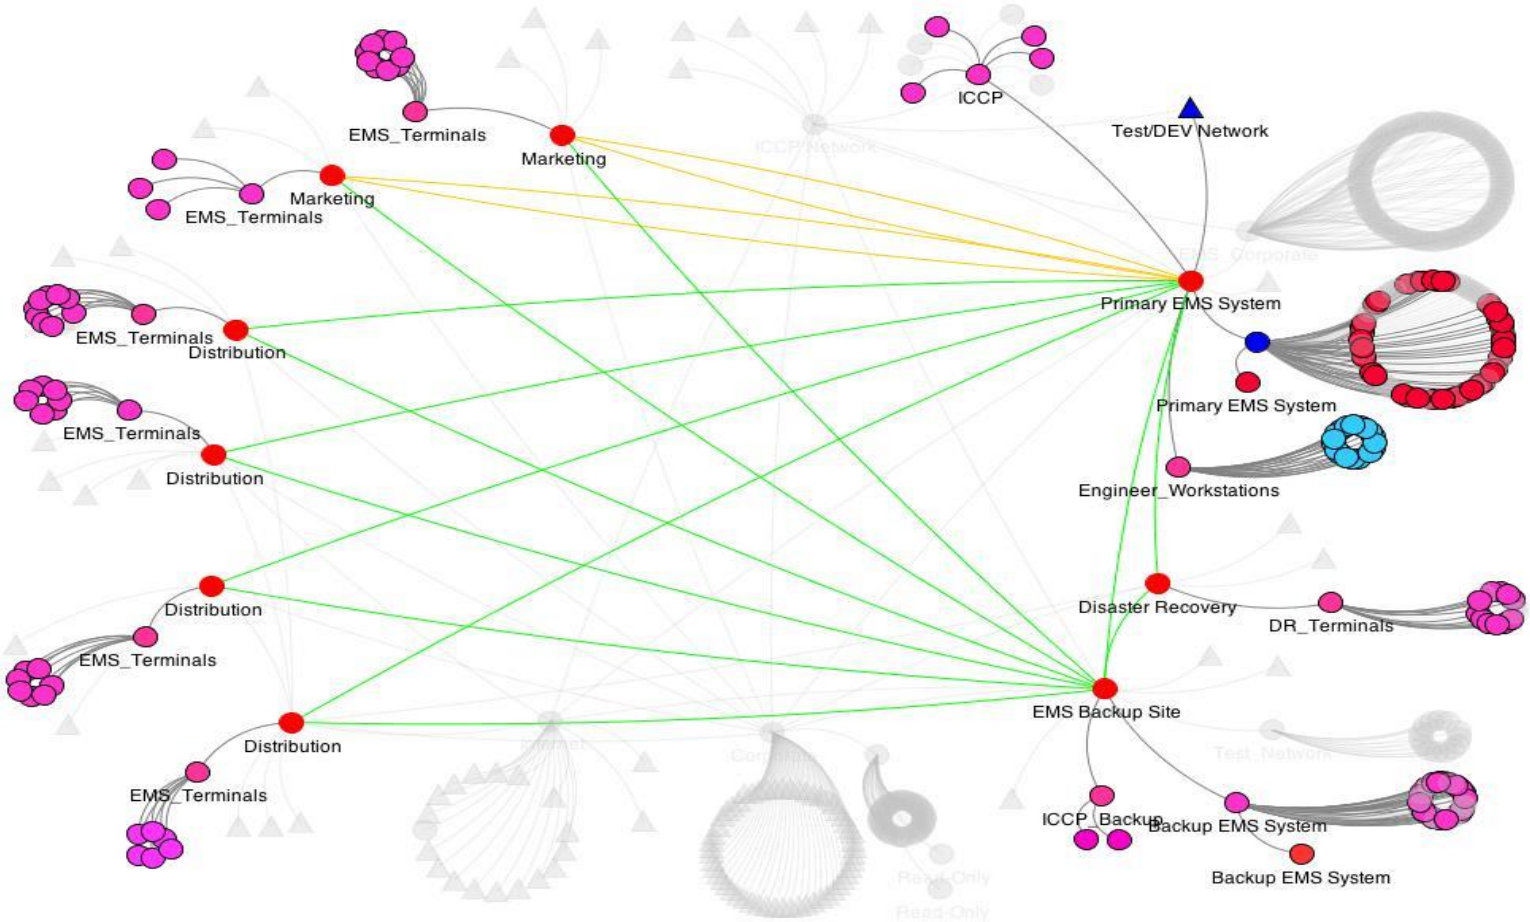
ICCP
Test/DEV Network
EMS_Terminals
Marketing
Marketing
EMS_Terminals
Primary EMS System
EMS_Terminals
Distribution
Primary EMS System
EMS_Terminals
Distribution
Engineer_Workstations
Distribution
Disaster Recovery
DR_Terminals
EMS_Terminals
EMS Backup Site
Distribution
EMS_Terminals
ICCP_Backup
Backup EMS System
Backup EMS System

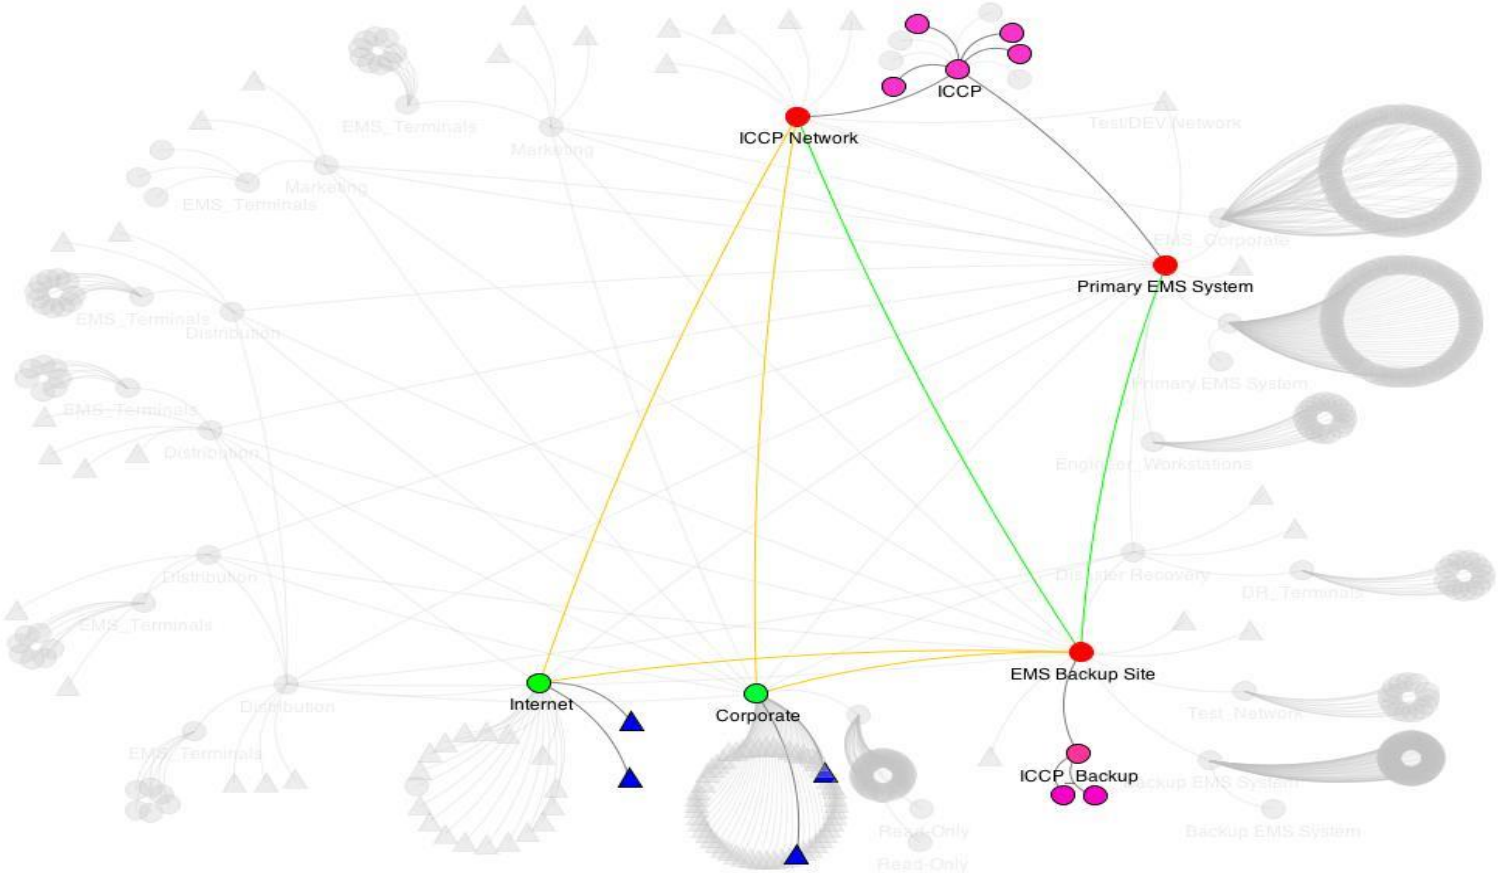
ICCP
ICCP Network
Primary EMS System
EMS Backup Site
Internet
Corporate
ICCP Backup

NetAPT for the auditor

- NetAPT provides a means for auditors to review security controls inside ESP's without having to directly access the ESP
- Time spent in review of firewall configurations can be significantly reduced
- Standardizing on a tool such as NetAPT would allow for consistent presentations of this sensitive data across utilities