The Emperor's New Cloud: An Analysis of the July 2009 RoK/ USA DDoS Attacks

Roland Dobbins
<rdobbins@arbor.net>
Solutions Architect
+66-83-266-6344 BKK mobile
+65-8396-3230 SIN mobile

“We cannot solve problems by using the same kind of thinking we used when we created them.”

- Albert Einstein

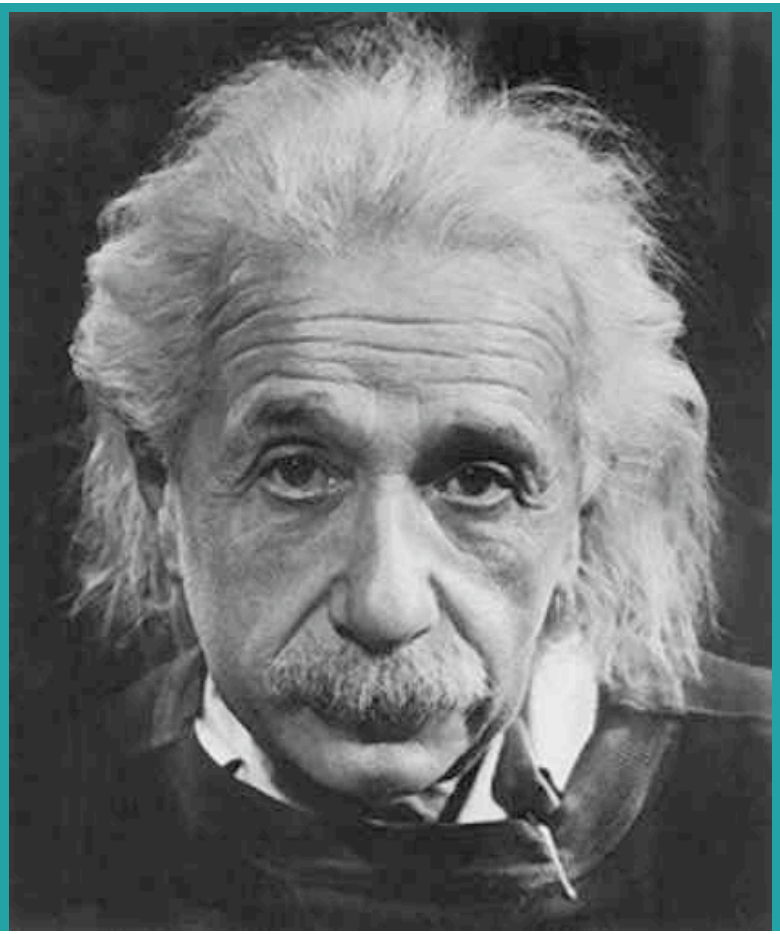

Introduction & Context

Evolution of Threats and Exploits

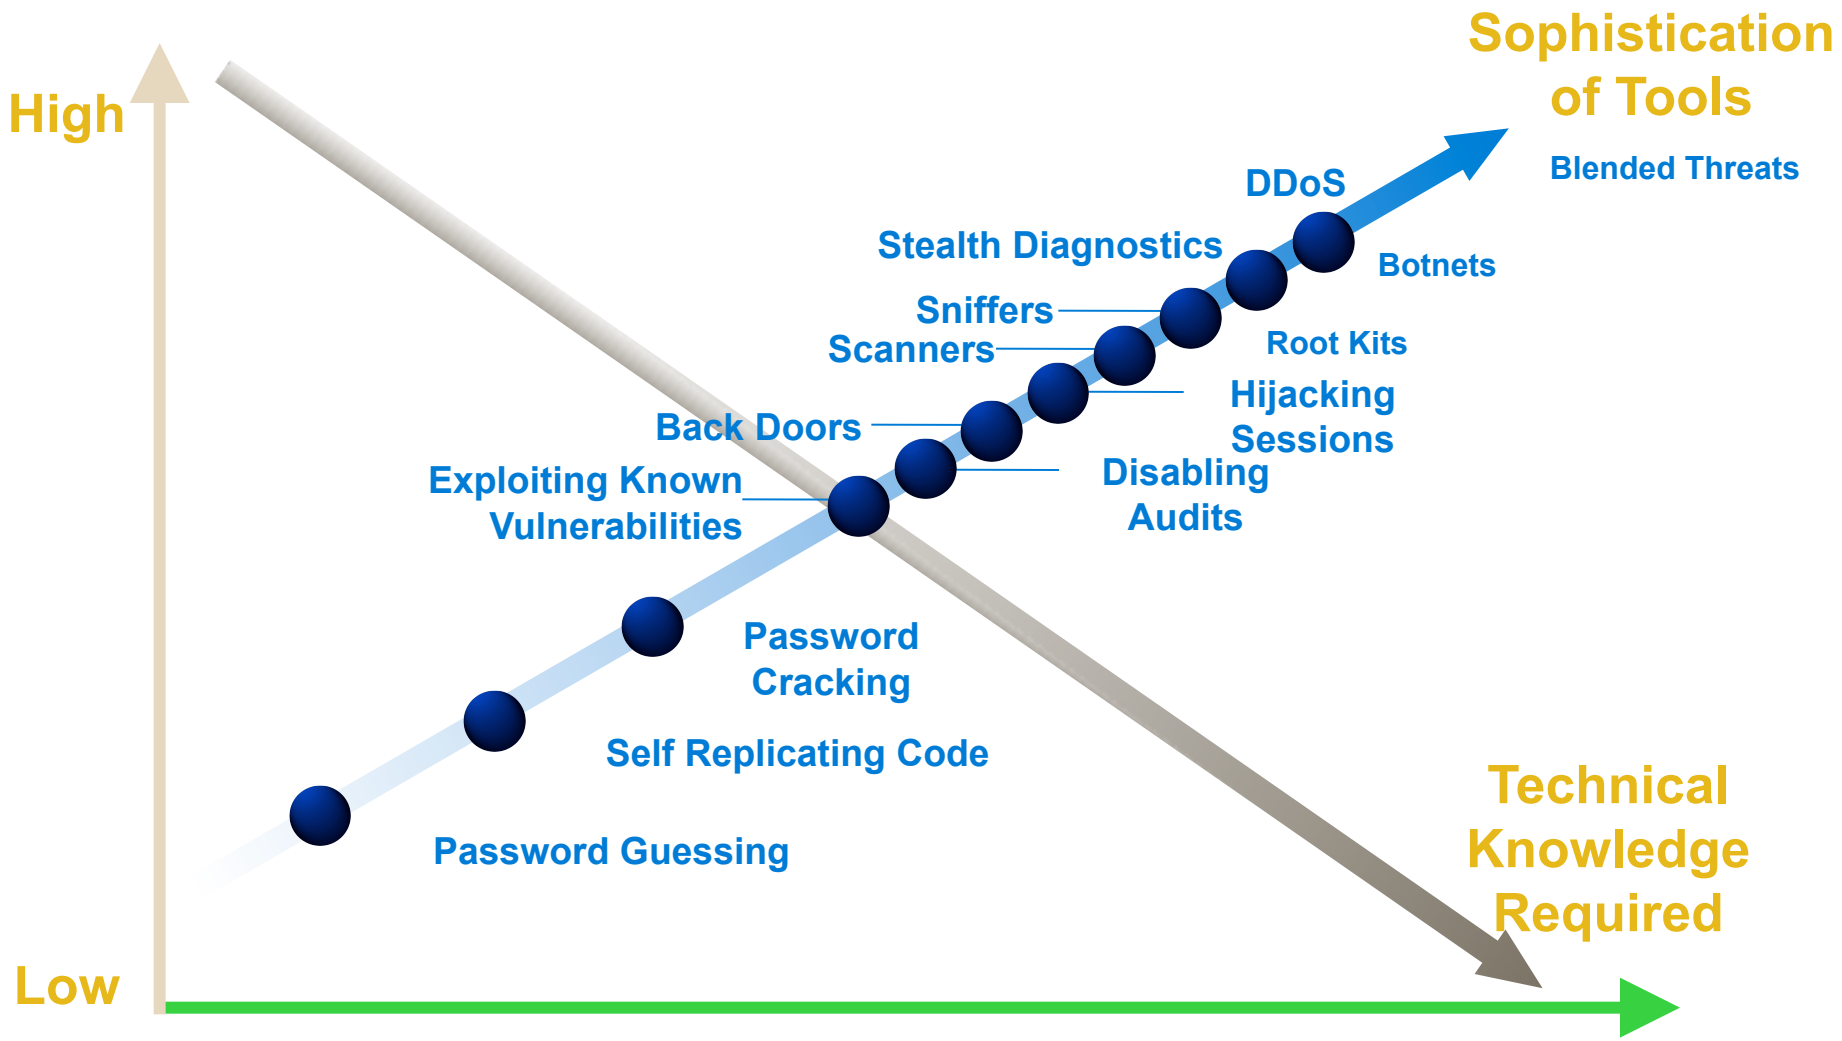

Botnets - The #1 Online Security Threat

Wikipedia on Botnets: **. . . a collection of compromised computers (called zombie computers) [or bots] running programs, usually referred to as worms, Trojan horses, or backdoors, under a common command and control infrastructure.**

Botnets are the prime enablers of all these activities:

- DDoS
- Extortion
- Advertising click-through fraud
- Fraudulent sales
- Identity theft and financial fraud (phishing, stealing info from PCs, etc.)
- Theft of goods/services
- Espionage/theft of information
- Spam-based stock-market manipulation

DDoS Attacks – A Fact of Life on the Internet

- DDoS attacks are taking place 24/7/365 – they're simply a fact of life on the Internet.
- Subjectively speaking, ~15% of DDoS attacks are financially motivated (mainly extortion), ~15% criminal retribution (phishers/ spammers vs. antispam, miscreants vs. miscreants); ~1% ideologically motivated; ~69% simple nihilism.
- Any organization, any site, any individual can be affected by DDoS, either as a direct target or via collateral damage.
- Outbound DDoS can be just as devastating to end-customers and SPs as inbound DDoS – botted hosts on broadband access networks, on enterprise networks, and within IDCs affect both the source networks and the targets.
- Situational awareness is key – what's happening in the news? What anniversaries are taking place this year/month/week/today?
- Miscreants attack one another with regularity – collateral damage!

DDoS Attacks – The #1 Security Threat to Cloud Computing!

- Discussions of cloud security priorities tend to focus on confidentiality of data, privacy, separation of application logic in a multi-tenanted cloud infrastructure.
- The cloud security elephant no one wants to discuss is DDoS - why?
- DDoS is the #1 security threat to the cloud model – DDoS shuts the cloud down. No cloud availability = no revenue for cloud providers!
- Most security researchers tend to ignore DDoS – why?
- Cloud providers don't use resiliency against DDoS as a selling point for their services – why?
- The reality is that we've all been dependent upon 'the cloud' for years – search engines, Web mail, IM, social networking, weblogs, etc. – and it's disruptive for ordinary users when these services are unavailable. What will it be like when even more core information infrastructure is dependent upon continuous Internet-wide availability? Netbooks, mobile apps/data?

The Emperor's New Cloud

- We're relying upon 25-year-old protocols designed for use in a laboratory environment and with little/no regard for security as the foundation of our global Internet infrastructure.
- Although there's a large body of work on operational security (opsec) and scalable Internet architectures, it's honored more in the breach than in actual deployments.
- Ongoing, pervasive disconnect between network architects, application architects, operational groups, security teams, management.
- Pollyannaish attitude towards security – 'Why would anyone attack *us*?'
- Lack of accountability – is anyone ever fired as a result of avoidable security incidents?
- Pervasiveness of security theater/security snake-oil.
- Inability/unwillingness to properly assess abstract threat models – a necessary psychological defense mechanism?

Analyzing the Attacks

The July RoK/USA DDoS Attacks – A Timeline

- First attacks against US targets seen on 5Jul09 APAC/4Jul09 USA – Independence Day holiday/long weekend for USA.
- Because it's a weekend, many are slow to notice attacks, even slower to coordinate responses.
- Attacks continue over the long weekend into the next week.
- Press/'blogosphere' coverage picks up on Monday, 6Jul09 USA.
- First attacks against RoK targets seen late on 7Jul09 – 8Jul09 is 15th anniversary of Kim Il-Sung's death, national day of mourning in North Korea.
- Confusion and lack of communications hinder responses in many cases.
- Botnet begins to self-destruct over the subsequent weekend.

RoK/USA DDoS Attack Traffic

- TCP/80 SYN-floods – most common DDoS attack type, directed against Web services on TCP/80.
- UDP/80 packet-floods – we see this from time to time, nonsensical, probably a default for some bots (miscreants often don't know much about networking).
- ICMP echo-reply floods – i.e., ping-floods. Common DDoS attack type.
- HTTP GET for '/' – common layer-7 attack, generally indicates no prior reconnaissance of target sites (i.e., no attempt to hit CGI scripts, pull down large files, etc.).
- HTTP GET for /china/dns? – no clue.
- Protocol 0 packet-floods – attackers sometimes use uncommon Internet protocols (i.e., something other than TCP/UDP/ICMP) because sites often fail to filter these protocols, packets make it through incomplete ACLs.

RoK/USA DDoS Attack Traffic (cont.)

- Relatively low observed pps/bps – ~25-50mb/sec, ~50-100kpps, in most instances. Largest single bps/pps recorded against a single target ~140mb/sec, ~500kpps.
- Unconfirmed reports of up to 2gb/sec – no data to support this.
- Some standard, some bogus HTTP header info in layer-7 attack components:

```
Mozilla/4.0 (compatible; MSIE 7.0; Windows NT 5.1; GTB6; .NET CLR 2.0.50727; .NET
CLR 3.0.4506.2152; .NET CLR 3.5.30729)

Mozilla/4.0 (compatible; MSIE 7.0; Windows NT 5.1; Trident/4.0; GTB6; .NET CLR
2.0.50727; .NET CLR 3.0.4506.2152; .NET CLR 3.5.30729)

Mozilla/5.0 (Windows; U; Windows NT 5.1; en-US; rv:1.8.1.20) Gecko/20081217 Firefox/
2.0.0.20 (.NET CLR 3.5.30729)

Mozilla/4.0 (compatible; MSIE 8.0; Windows NT 5.1; Trident/4.0; GTB6; .NET CLR
2.0.50727; .NET CLR 3.0.4506.2152; .NET CLR 3.5.30729
```

RoK/USA DDoS Attack Traffic (cont.)

- Additional header components:

```
Accept: image/gif, image/x-xbitmap, image/jpeg, image/pjpeg, application/x-
shockwave-flash, application/vnd.ms-excel, application/vnd.ms-powerpoint,
application/msword, application/x- ms-application, application/x-ms-xbap,
application/vnd.ms- xpsdocument, application/xaml+xml, */*

Accept-Language: ko

UA-CPU: x86

Accept-Encoding: gzip, deflate

Content-Length: 0

Connection: Keep-Alive
```

RoK/USA DDoS Attack Methodology

- Apart from initial bot target-list seeding, attack methodology seemed largely based upon hands-on manual guidance of bots.
- Attacker shifted targets seemingly in response to successful defense – often hours after the attack had been mitigated at a particular site. This indicates manual monitoring of attack impact, only semi-active targeting – not real-time.
- Attacker also seemed to shift targets semi-randomly.
- Attacker never varied attack traffic mix – implies inexperience, lack of technical acumen.
- Attacker apparently didn't perform prior reconnaissance, didn't customize attack traffic/mechanisms for high impact – again, implies inexperience and lack of technical acumen.
- Attacker seemed focused on the conjunction of USA Independence Day holiday/15th anniversary of death of Kim Il-sung. Attacks initially launched via timers configured in bots.
- Authoritative DNS server outages for some target sites apparently collateral damage due to non-scalable DNS infrastructure, high rates of attacking bot & legitimate user DNS queries.

RoK/USA DDoS Botnet Details

- Largely derived from older MyDoom-A/B variants – codebase 5-6 years old. Detected by modern antivirus as MyDoom.
- Multiple C&C servers in several countries outside of RoK. C&C accomplished via static config files on C&C servers, bots checked in 1/hour. Targets updated 1/day during most of the attacks, 2/day on one occasion.
- 95% of botted hosts within RoK; botted hosts outside RoK apparently belonged to RoK nationals living/working abroad.
- ~130K bots verified; unverified reports of ~200K bots.
- Malware compromise vector used to build botnet RoK-specific. Semi-social-engineering, no technical innovation, somewhat novel approach.
- Bots made use of local system clock for timed events – attack start, self-destruction.
- Bots self-destructed starting on 10Jul09 according to system clock – deleted/compressed files, overwrote MBR. Possible for informed users to stop self-destruction by resetting system clock, etc.

RoK/USA DDoS Attack Impact

- Some targeted USA governmental sites experienced total and partial service outages, recovered relatively slowly, mostly after the long holiday weekend. USA ASPs/commercial & some governmental sites largely unaffected.
- Many targeted RoK commercial and governmental sites experienced total and partial service outages; most recovered slowly. Many targeted RoK sites continued to experience outages until the onset of botnet self-destruction.
- Due to more clueful selection of RoK attack targets, much more impact on the day-to-day lives of RoK individuals.
- No reported broadband outages/impact due to botnet DDoS traffic – probably due to relatively low pps/bps, lack of DNS attack component.
- RoK bore the brunt of meaningful impact – USA impact largely marginal due to uninformed attack target selection, generally greater level of preparedness. Korean press reported one auction site suffered as much as $6M USD in lost sales.
- Overall impact blunted due to ineptitude and lack of technical acumen of the attacker.
- These were stupid, unoriginal attacks; no innovation. Observed impact largely due to ***unpreparedness of defenders***.

Successful Defense – What Worked?

Practice(s) Makes Perfect!

- Organizations with strong communications plans, contacts for peers/ upstreams, opsec teams were able to react much more quickly and effectively.
- Organizations with strong, scalable architectures saw hardly any impact at all. Best Current Practices (BCPs) work!
- Organizations with detection/classification capabilities saw the attack traffic, characterized it quickly, understood what was happening.
- Organizations with source-based remotely-triggered blackholing capability (S/RTBH) used this tool effectively to block attacking bots.
- Organizations with intelligent DDoS mitigation systems (IDMS) used these tools effectively to block attacking traffic on a more granular level.
- Online mitigation communities effective in coordinating responses.

Unsuccessful Defense – What Didn't Work?

Ignorance isn't bliss!

- Organizations without strong communications plans, contacts for peers/ upstreams, opsec teams were unable to react quickly and effectively.
- Organizations without strong, scalable architectures suffered continuous outages during the attacks.
- Organizations without detection/classification couldn't see the attack traffic (or even that an attack was taking place), had no idea what was happening.
- Organizations without mitigation/reaction tools such as S/RTBH and IDMS had no scalable, operationally feasible way to react.
- Organizations with firewalls and IDS/'IPS' inline in front of their servers went down quickly and stayed down. Same for load-balancers.
- Organizations which weren't members of online mitigation communities and weren't tuned into the global security community suffered extended outages.

How Can We Defend Against DDoS Attacks?

Pervasive Security in an Age of Distrust

- **Security is the heart of internetworking's future; we have moved from an Internet of implicit trust to an Internet of pervasive distrust**
- **Network/application design = security, security = network/application design**
- **We can no longer differentiate networking & applications from security, they must be intertwined**

  What is security? QoS? Routing? DNS? Web 2.0?
- **No packet can be trusted; all packets must earn that trust through a network device's ability to inspect and enforce policy**

Three Security Characteristics

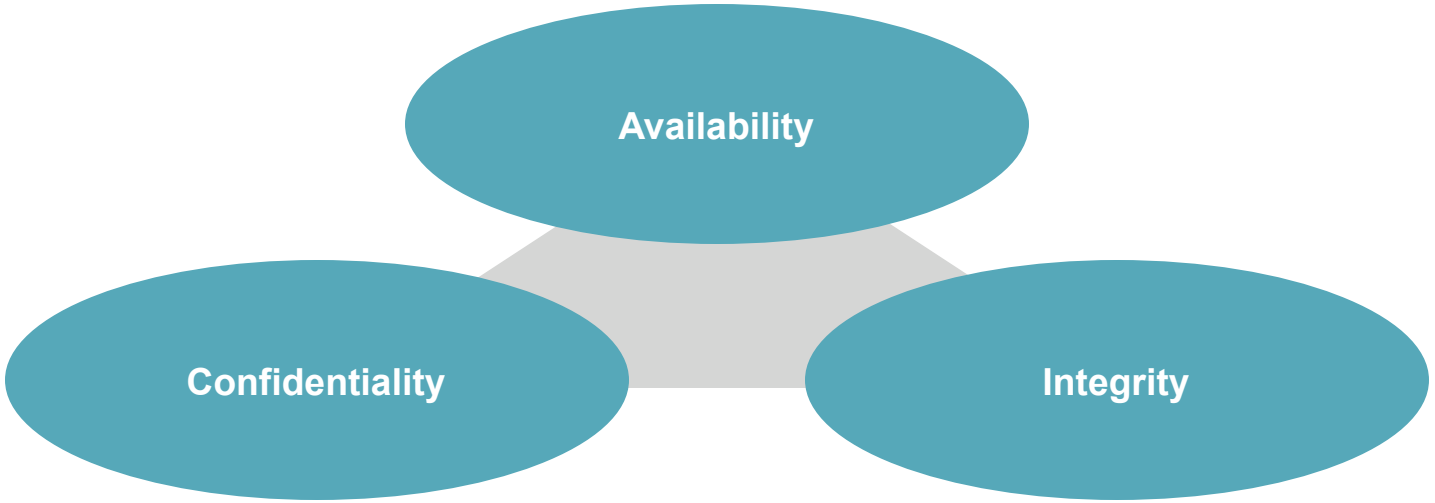

- The goal of security is to maintain these three characteristics

Three Security Characteristics

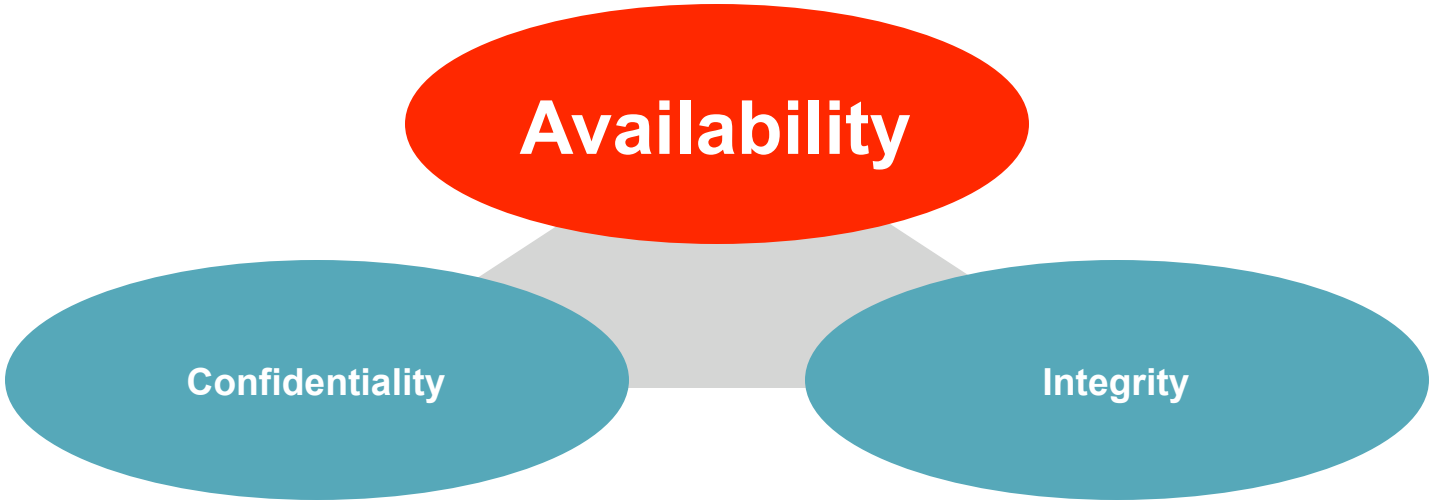

- Primary goal of infrastructure security is maintaining availability

Network/Application Availability: Protect the Infrastructure

- Security is the heart of internetworking's future; we have moved from an Internet of implicit trust to an Internet of pervasive distrust
- No packet can be trusted; all packets must earn that trust through a network device's ability to inspect and enforce policy
- Protecting the infrastructure is the most fundamental security requirement
- Infrastructure protection should be included in all high availability designs
- A secure infrastructure forms the foundation for continuous service delivery

Six Phases of Incident Response

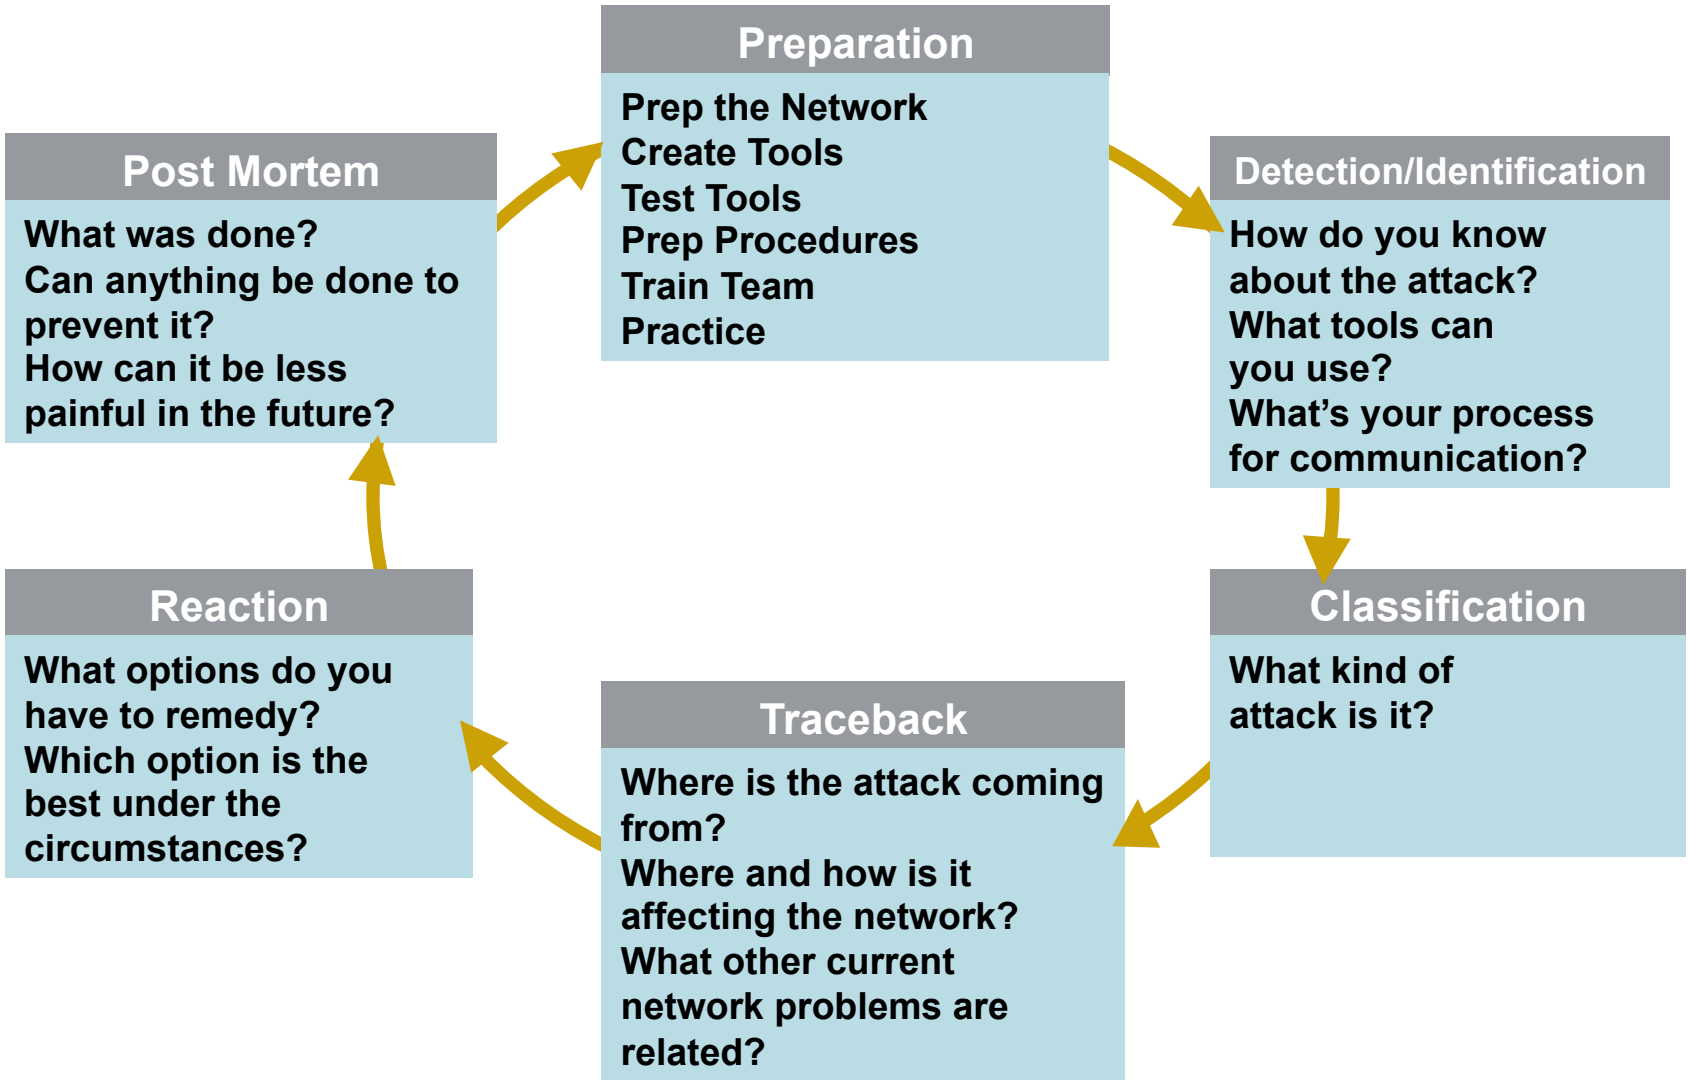

Are You Pushing the Envelope?

Know Your Equipment and Infrastructure:

- Know the performance envelope of all your equipment (routers, switches, servers, etc.). You need to know what your equipment is really capable of doing.
- Know the capabilities of your network. If possible, test it. Surprises are not amusing during a security incident.
- PPS vs. BPS and, how enabling features impacts them

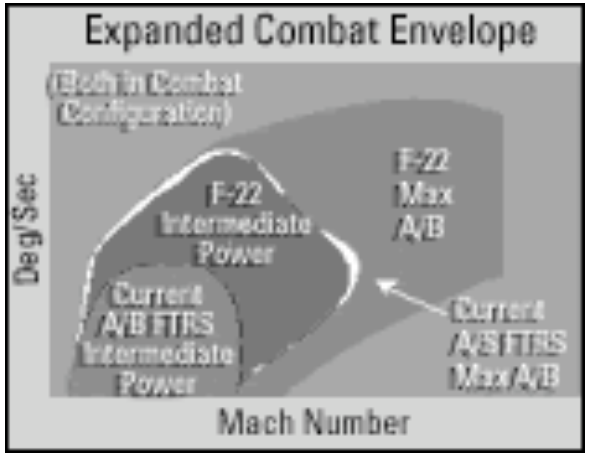

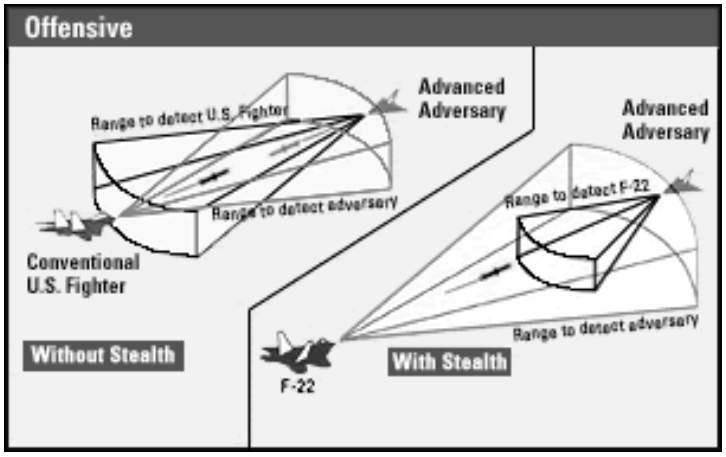

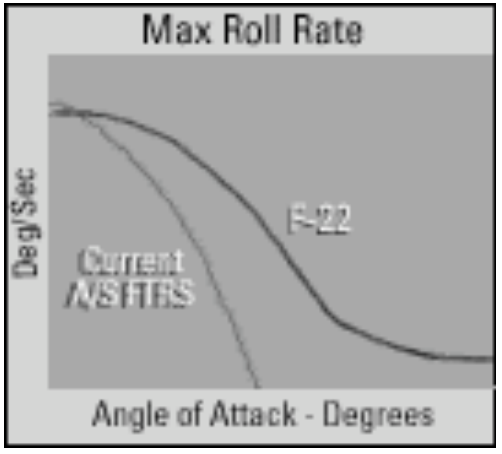

Architecture

The Right Tools for the Right Job

Infrastructure Best Current Practices (BCPs)

- Interface ACLs (iACLs) should be employed at the relevant network edges (peering/transit, customer aggregation edge, etc.) to protect the network infrastructure itself; additional service-specific sections should be used to restrict traffic destined for Internet-facing servers to the ports and protocols associated with the services and applications on those servers.
- The use of IP protocol 0 in this attack is notable as a common mechanism used by attackers to bypass ACLs that only contain policy statements relating to common protocols such as TCP, UDP, and ICMP; there are 254 valid Internet protocols, and irrelevant protocols should be filtered at the edges via ACLs.
- Additional network infrastructure BCPs such as control- and management-plane self protection mechanisms (rACL, CoPP, GTSM, MD5 keying, et. al.) should also be deployed.
- All network infrastructure devices should be accessible only via designated management hosts, and this access should be facilitated via a dedicated out-of-band (OOB) management network. During high-impact DDoS attacks, a dedicated management network ensures that devices can be managed irrespective of conditions on the production network, and also ensures that vital mechanisms such as flow telemetry and SNMP are uninterrupted, which assures continuing visibility into attack traffic during an incident

Infrastructure BCPs (cont.)

- Flow telemetry such as Cisco NetFlow, Juniper cflowd, and sFlow should be enabled at all network edges, and exported into a collection/analysis system.
- Source-based remotely-triggered blackholing (S/RTBH) is a powerful reaction technique which allows tens or even hundreds of thousands of attacking source IPs (classified via flow analysis, logfiles, etc.) to be rapidly blackholed based upon their source addresses. S/RTBH leverages BGP as a control-plane mechanism to instantaneously signal edge devices to start dropping attack traffic.
- Intelligent DDoS mitigation systems (IDMS) should be deployed in topologically-suitable cleaning centers in order to protect servers/services/applications. They should be emplaced northbound of load-balancers; if an organization insists on placing firewalls and IDS/'IPS' inline in front of servers, protect these stateful DDoS chokepoints and everything behind them!
- Do not place firewalls and IDS/'IPS' in front of servers – they provide no security value whatsoever in server environments where every incoming connection is by definition unsolicited. They are DDoS chokepoints, and degrade the operational security posture of the network and applications.
- Policy should be enforced by stateless ACLs in hardware-based routers/ switches!

Host Best Current Practices (BCPs)

- Public-facing servers should be configured in a hardened manner, with unnecessary services disabled, OOB management access, service-specific configuration hardening, IP stack tuning, and other relevant mechanisms.
- Stateless on-server filtering via tcpwrappers is a useful policy-enforcement mechanism; for Web servers, Apache modules such as mod_security and mod_evasive bring additional capabilities.
- The deployment of stateful firewalls or other inspection devices such as IDS/'IPS' in front of Internet-facing servers is contraindicated; as each incoming connection to Internet-facing servers is by definition unsolicited, the stateful inspection adds nothing to the security posture of the servers, and serves to weaken their ability to withstand DDoS traffic due to the limited state-table size of even the largest/fastest firewalls and IDS/IPS on the market today.

  During this particular attack and during many other attacks, Web application firewalls in front of targeted servers were observed to fail while receiving relatively low amounts of attack traffic, thereby enabling the DDoS to succeed in making the servers unavailable with little effort on the part of the attacker

Host BCPs (cont.)

- Load-balancers also instantiate state which renders the real servers behind the load- balancers more vulnerable to DDoS; during this attack, load-balancers were observed to fail due to state exhaustion as a result of the attack traffic. S/RTBH, reverse-proxy caches, & IDMS should be utilized to protect the load-balancer and the real servers behind it.

- DNS infrastructure should be deployed in a modular, bulkheaded architecture, with separation of functions such as authoritative servers, internal resolvers, external resolvers, caching-only resolvers, etc., and should be scaled appropriately by employing techniques such as IPv4 anycast addressing. S/RTBH & IDMS should be employed to protect the DNS from deliberate attack and/or collateral damage.

  During this attack, it does not appear that DNS was directly targeted; however, DNS lookups for targeted domains were observed to intermittently fail, probably as a result of large amounts of queries for sites in these domains by both legitimate users attempting to access them repeatedly as well as by botted hosts resolving the target host IPs as part of the HTTP GET attack component.

  This is a classic example of collateral damage to a vital infrastructure service.

A Logically-Separated, Bulkheaded DNS Architecture

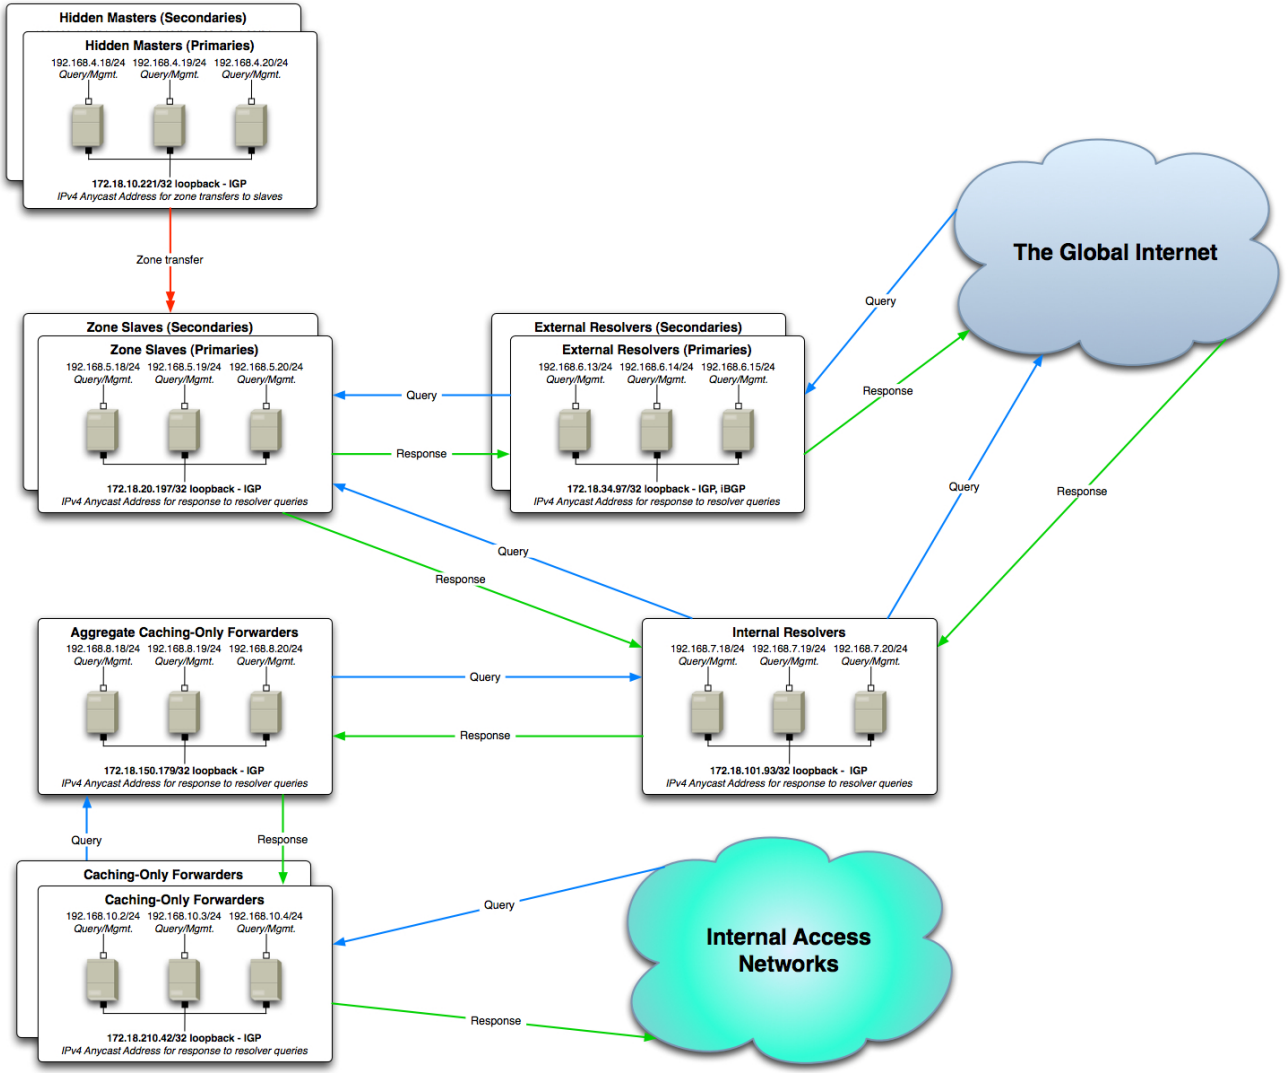

The Right People for the Right Job

OPSEC Team Skill Requirements

- **The OPSEC Team needs to know ….**
  - Everything a Backbone Engineer knows
  - Everything a Network Management Engineer knows
  - Everything a sysadmin/webmaster knows
  - Everything an email postmaster knows
  - Everything a DNS/DHCP/Addressing Engineer knows
  - Everything a CERT Engineer knows
  - Everything an Enterprise Infosec specialist knows

In essence, you're looking for super-engineers who are hybrid Backbone/Security Engineers.

Conclusions

What Have We Learned?

- Plenty of tools, techniques, and BCPs exist which, if deployed by targeted organizations, would've rendered this attack completely ineffective.
- We seem to repeat the mistakes of the past over and over again – siloed organizations, lack of communications, lack of strategy.
- Organizations which plan and prepare ahead of time recover quickly and are resilient in the face of attack.
- The RoK/USA DDoS attacks of July 2009 were run-of-the-mill, ordinary, rather inept attacks; they were effective against many of the targets due solely to the lack of planning and operational readiness of the target organizations.
- The same was true of the Estonian DDoS attacks of 2007, and the Russia/Georgia/Azerbaijan DDoS attacks of 2008 – small, simple attacks which achieved disproportionate impact due to the unpreparedness of the defenders.
- We are building the cloud on a foundation of sand.

Are We Doomed?

- No! Deploying the existing, well-known tools/techniques/BCPs results in a vastly improved security posture.
- The challenge is to educate and empower architects and operators to do so.
- Much, much more education is needed at the BDM and TDM levels; we're doing this now, together here today.
- The cloud has the potential to bring a much higher degree of resilience and operational security to ordinary users – assuming the cloud providers Do The Right Thing with regards to tools/techinques/BCPs.
- We must look beyond IPv4 and IPv6 - which solves none of the security problems associated with IPv4, and introduces challenges of its own – and begin work on the next-generation transports and application-layer protocols which incorporate the security lessons we've learned over the last couple of decades.
- Automation is a Good Thing, but it's no substitute for resilient architecture, insightful planning, and plain old elbow-grease – top-notch opsec personnel are more important now than ever before!

Q&A

Special thanks to Jose Nazario & Danny McPherson of Arbor Networks for attack data and input into this presentation!

Thank You

Roland Dobbins <rdobbins@arbor.net>
Solutions Architect
+66-83-266-6344 BKK mobile
+65-8396-3230 SIN mobile