State of Danger:
Eliminating Excessive State in Network, Application, & Services Architectures as a DDoS Defense Strategy

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

Introduction & Context

DDoS Background

What is a Distributed Denial of Service attack?

- An attempt to consume finite resources, exploit weaknesses in software design or implementation, or exploit lack of infrastructure capacity
- Targets the *availability* and utility of computing and network resources
- Attacks are almost always *distributed* for even more significant effect – i.e., *DDoS*
- The *collateral damage* caused by an attack can be as bad, if not worse, than the attack itself
- DDoS attacks affect *availability*! No availability, *no applications/services/data/Internet*! No *revenue*!
- DDoS attacks are attacks against *capacity* and/or ***state!***

Three Security Characteristics

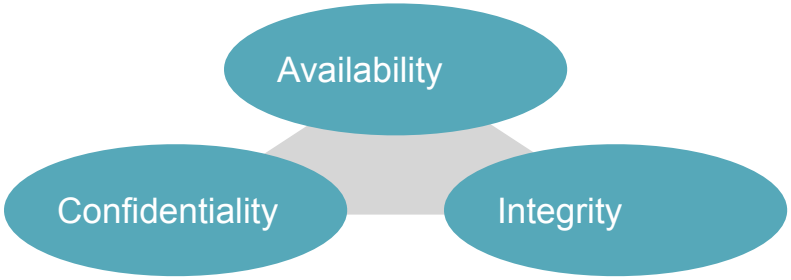

- The goal of security is to maintain these three characteristics

Three Security Characteristics

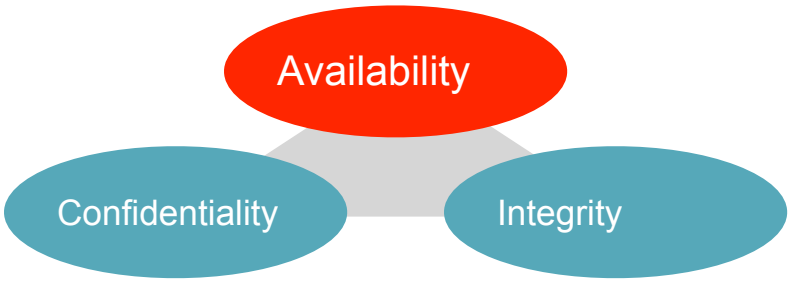

- Primary goal of DDoS defense is maintaining availability

State Exhaustion is the 'Silent Killer' of the Internet

- Most people tend to think about DDoS - if they think about it at all - in terms of bandwidth - i.e., bits/sec.
- In most (not all) volumetric attacks, throughput - i.e., packets/sec - is generally more important.
- In many cases, state exhaustion - overwhelming the ability of a device which makes packet forwarding decisions at least in part by tracking connection status - is an even more important factor.
- There's lots of unnecessary state on the Internet today, and it seems as if the problem is only getting worse!

State Reduction in the Internet Data Center (IDC)

The State of State in the IDC

- For ordinary users, the network doesn't matter - what matters is the applications, services, and data they need in order to achieve their goals (run business applications, communicate via VoIP, play BF3, et. al.)
- Unfortunately, many (most?) Internet-facing applications/services/data repositories are designed and deployed with fragile, brittle, non-scalable architectures.
- In particular, unnecessary and avoidable state is a big contributor to said fragility, brittleness, non-scalability.
- State exhaustion is a huge DDoS vector - whether or not attackers realize that's what they're accomplishing!
- Lack of cross-functional skillsets and inadequate architectural guidance are key contributing factors.

4AM Call - “Help! Our *entire* IDC is down!”

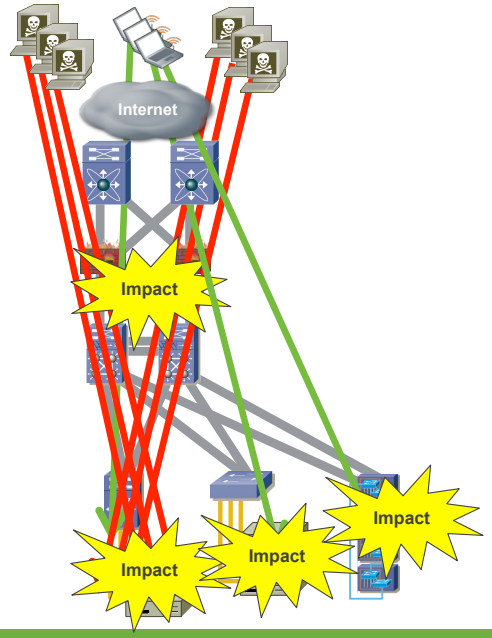

Segregate Traffic for Customers Who Insist Upon Stateful Firewalling - Limit Collateral Damage!

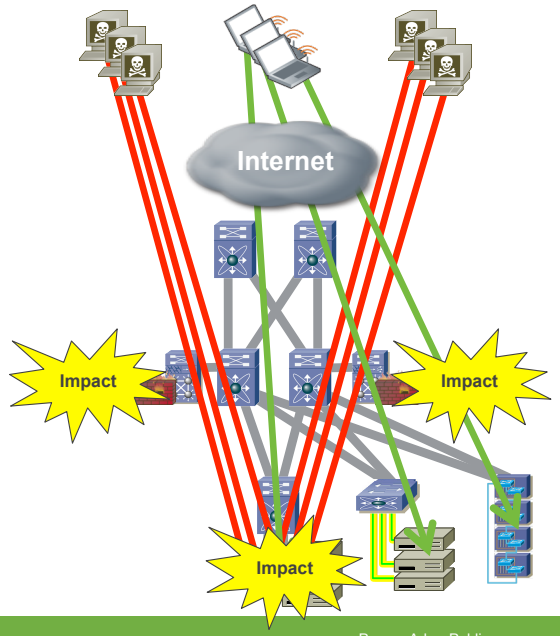

Stateful Firewalls in Front of Servers Considered Harmful!

- Why deploy a stateful firewall in front of servers, where every incoming connection is unsolicited, and therefore there is no state to inspect?!
- Policy enforcement can and should be accomplished via stateless ACLs in hardware-based routers and layer-3 switches capable of handling mpps!
- The 'inspectors' in stateful firewalls make things even worse - and they constitute a vastly expanded attack surface!
- In many (most?) cases, stateful firewalls are deployed as much due to organizational silioing/politics as to lack of technical acumen.
- AAA mechanisms in modern routers/switches can be used to allow appropriate security team access!
- If stateful firewalls cannot be immediately removed from the architecture, they must be protected against DDoS via S/RTBH, flowspec, IDMS, et. al., just like servers!

Arbor 6th Annual Worldwide Infrastructure Security Report - Stateful Firewall & IPS Failure Under DDoS

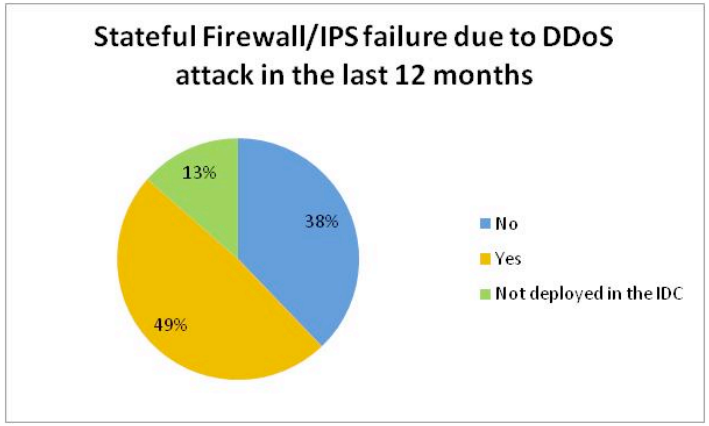

- **Nearly half of all respondents have experienced a failure of their firewalls or IPS due to DDoS attack!**

‘IPS’ Devices Carry Even More State!

- ‘IPS’ devices suffer from the same state-exhaustion issues as stateful firewalls - but even more so, as they typically try to hold multiple packets in memory simultaneously in an attempt to detect packet-level exploits.
- Attempted exploitation and compromise are table stakes for being on the Internet. Someone (or something) is *always* trying to hax0r you!
- The only way to secure servers/applications/services against exploitation and compromise is via secure architectural, coding, and maintenance (i.e., patching) BCPs.
- Why place an ‘IPS’ device on the Internet - after all, do you still have your email client set to alert you to incoming mail? ;>
- If ‘IPS’ devices cannot be immediately removed from the architecture, they must be protected against DDoS via S/RTBH, flowspec, IDMS, et. al., just like servers!

Load-Balancers Are Stateful Devices, Too!

- Load-balancers suffer the same challenges as stateful firewalls with regards to state exhaustion - in many cases, load-balancers go down under trivial DDoS attacks.
- There are many different mechanisms available to perform load-balancing other than dedicated load-balancing devices - Pen, Pound, LVS, Balance, Apache Traffic Server, mod_proxy_balancer, etc.
- Load-balancers must be protected against DDoS - stateless ACLs for policy enforcement, S/RTBH, flowspec, IDMS, and so forth.
- Fronting load-balancers with reverse proxy-caches is an architectural BCP (more on this later).

A Salient Comment on PCI/DSS.

"PCI should be more risk-based with more options, and less that is proscriptive; it's both too proscriptive and too vague at the same time."

-- Michael Barrett, PayPal CISO

PCI/DSS Compliance Does *Not* Require Stateful 'Application Firewalls'!

- Contrary to popular belief (and vendor propaganda), PCI/DSS compliance for organizations/sites which handle credit card payments does *not* require stateful 'application firewalls' to be placed in front of Web servers.
- On-node, integrated solutions such as mod_proxy (free!) and URLScan (free!) meet all the PCI/DSS requirements for 'application firewalls' - and they aren't stateful network DDoS chokepoints which will bring down your entire application stack!
- If your PCI/DSS auditor disagrees, a bit of education generally does the trick.
- If not - find another PCI/DSS auditor! ;>

State Reduction in Mobile Wireless Networks

Legacy Aspects of Mobile Wireless Architectures

- Until recently, most mobile wireless networks were designed and built with ‘minutes’ in mind - data was an afterthought, and the emphasis was on highly skilled/specialized folks on the ‘minutes’ side of things, rather than TCP/IP.
- With the rise of iDevices, many mobile wireless have essentially become ‘accidental ISPs’.
- Because of the technical emphasis on ‘minutes’, many BCPs were not implemented; many mobile wireless networks were designed in much the same fashion as (brittle, fragile, non-scalable) enterprise networks, containing excessive state in the form of NAT and stateful firewalling.
- Many mobile wireless networks suffer from availability issues directly related to outbound/crossbound botnet activities, including DDoS, as a result.

4AM Call - “Help! Our *entire* 3G network is down!”

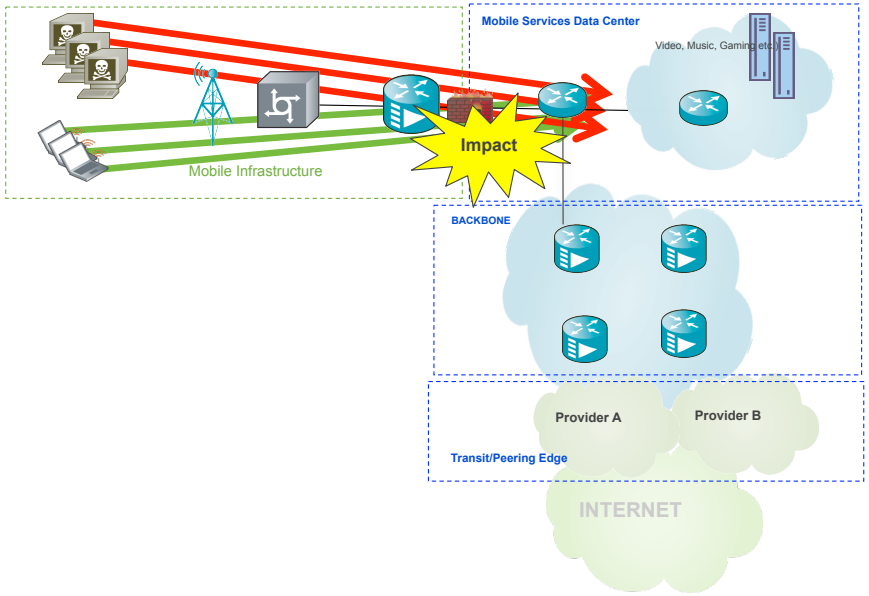

4AM Call - “Help! Our *entire* 3G network is down!”

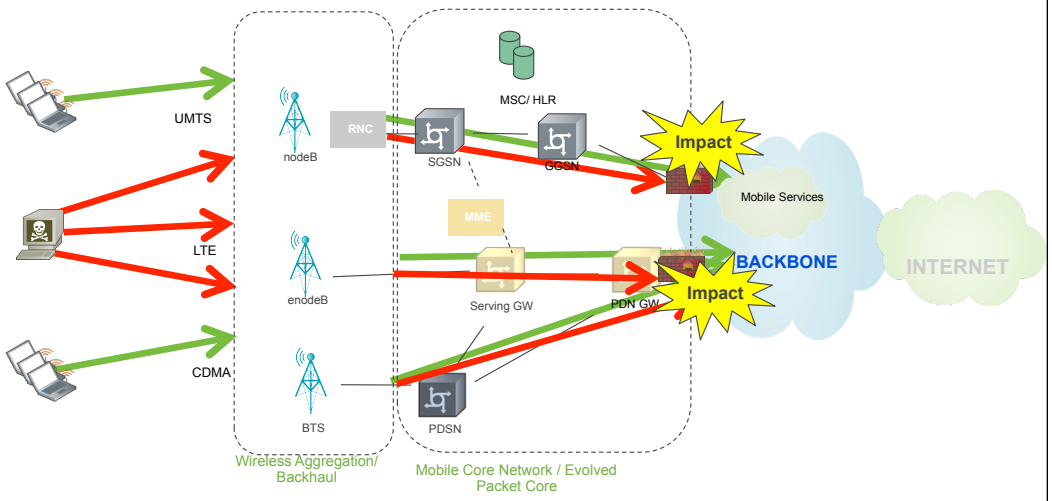

Stateful Firewalls (and NAT) in Mobile Wireless Networks Considered Harmful!

- Stateful firewalls are not deployed in the data plane of (almost all) wireless broadband networks for a reason!
- NAT isn't performed above the CPE level in (almost all) wireless broadband networks for a reason (more on this later)!
- It is possible to design mobile wireless data networks today without using NAT.
- It is possible to use stateless ACLs in hardware-based routers and layer-3 switches in order to keep almost all externally-originating scanning activity from 'waking up' mobile subscriber nodes.
- If stateful firewalls and/or NAT devices can't be immediately removed from mobile wireless networks, those devices must be protected to the degree possible against DDoS attack via S/RTBH, flowspec, IDMS, quarantine systems, et. al.

State Reduction in Application Delivery Architectures

Minimize/Eliminate State on the Front-End!

- Applications should be designed in such a way that all application state is handled at layer-7 - there should be no stateful tracking performed based upon TCP/IP semantics. This allows horizontal scalability of the front-end and middle-tier servers (database/ datastore architectures are beyond the scope of this presentation).
- Reverse-proxy caches such as Squid, Varnish, NGINX, HAProxy, mod_proxy, et. al. should be deployed for HTTP-based applications. Packets from outside your network should never be allowed to touch your actual front-end servers, load-balancers, etc. WCCP is a Good Thing, too!
- For other applications, make use of generic front-end reverse-proxies as much as possible; use custom code as necessary. Do not let packets from outside your network touch your real front-end servers and/or load-balancers!
- Reverse-proxy farms must be protected from DDoS via S/RTBH, flowspec, IDMS, et. al.
- Make use of memcached, etc. as appropriate - again, no packets from outside!

IPv6 - Bringing Mobile Wireless-Style Stateful DDoS Chokepoints to a Wireline Network Near You!

In the Medium Term, IPv6 Migration Will Bring More State, Not Less.

- Myth - IPv6 means no NAT.
- Reality - with IPv4 address exhaustion looming, Carrier Grade NATs (CGNs) are being deployed on SP wireline networks.
- 6-to-4 gateways are stateful devices with the same issues as those surrounding NAT devices. 6-to-4 gateways were being deliberately DDoSed back in 2004.
- Many of the performance/latency issues associated with mobile wireless networks will make their way into wireline networks as a result.
- These stateful devices must be protected to the degree possible against DDoS attack via S/RTBH, flowspec, IDMS, quarantine systems, et. al.

Huge Amounts of Excessive, Harmful State Are the ‘Elephant in the Room’ of the Transition to IPv6!

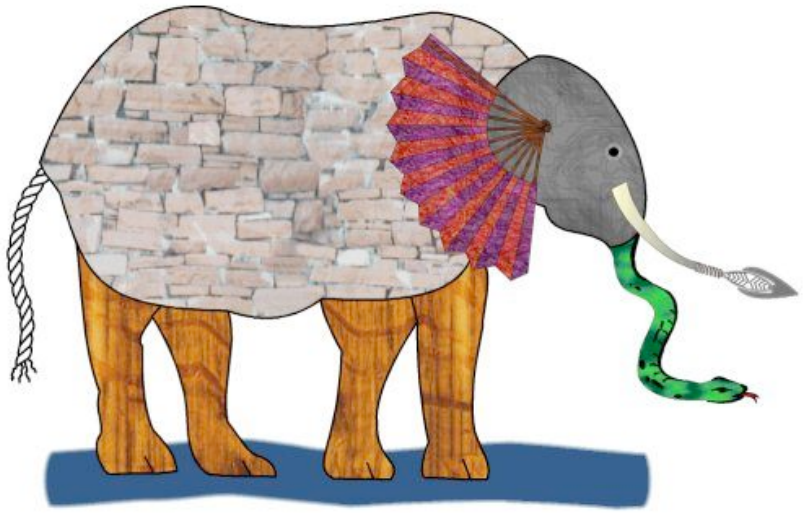

Are We Moving Towards a Less Resilient Internet as a Result of IPv6 Migration & Related Trends?

Conclusions

Conclusions

- Excessive, unnecessary state is a barrier to scalability and lowers resilience to DDoS attacks.
- Many DDoS attacks are successful due solely to state exhaustion of stateful firewalls, 'IPS' devices, load-balancers, etc.
- Stateful firewalls should not be placed in front of servers; if they can't be removed, they must be protected against DDoS attacks.
- IPS devices should not be placed in front of servers; if they can't be removed they must also be protected against DDoS attacks.
- Ditto for load-balancers.
- Policy enforcement should be implemented via stateless ACLs in hardware-based routers/layer-3 switches
- Applications and their delivery infrastructures should be designed in such a way as to minimize unnecessary state.
- The transition to IPv6 is going to result in more NAT, not less, and more stateful devices such as 6-to-4 gateways, not fewer.
- Education and opex are the keys to maintaining availability!

Q&A

ARBOR NETWORKS

Thank You!

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