China in the Carna Botnet

Telnet's threat to an emerging superpower

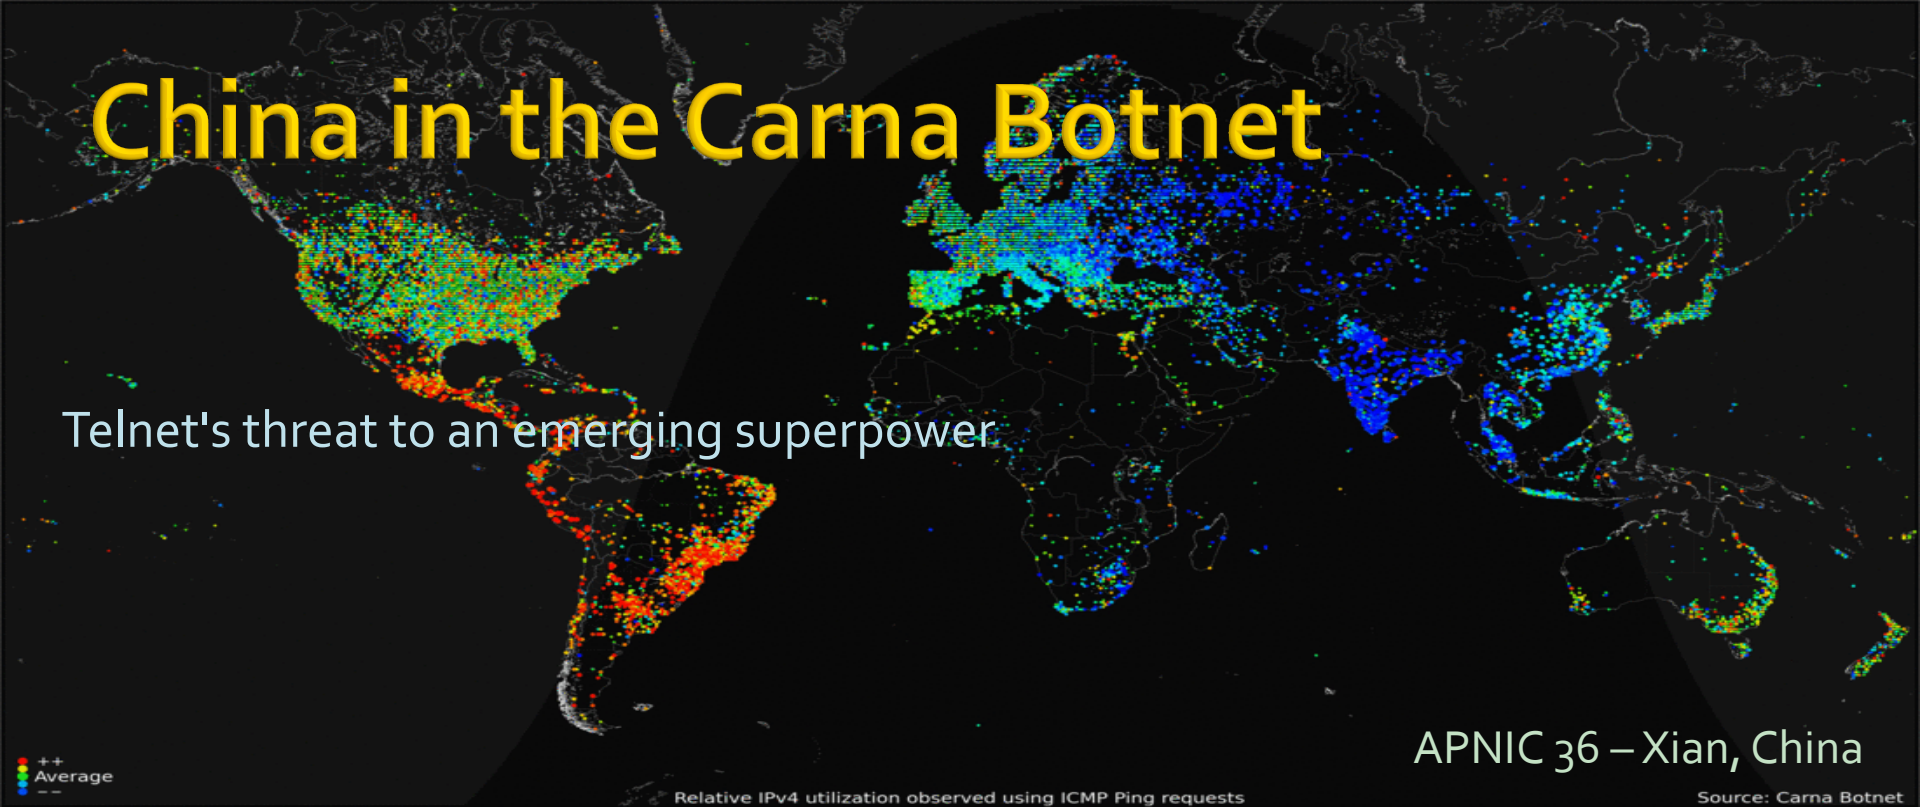

APNIC 36 – Xian, China

Parth Shukla
Information Security Analyst,
Australian Computer Emergency
Response Team

pparth@auscert.org.au
twitter.com/pparth

Why should you care?

- This research shows that there is a high ratio of easily vulnerable devices in China – which highlights a major security concern.
- Malicious agents can use this to take root control of vulnerable devices
- Can have a serious impact on Chinese Economy because root shell access means:
  - Sniff all passing network traffic, and/or
  - Modify passing network traffic, and/or
  - Shutdown/reboot devices at will, and/or
  - Use it to relay illegal traffic (such as child porn) and/or
  - Perform cyber attacks on other countries or companies and China or innocent companies will be blamed
  - Network operators will have their network unnecessarily clogged

AusCERT

What is the Carna Botnet?

- Millions of devices that were compromised for use in conducting the "Internet Census 2012" by an anonymous researcher
- 70% of these devices were either too small, didn't run Linux or were somehow limited (e.g. no "ifconfig")
  - Traceroutes of **some** of these devices are part of the public torrent
- ≈1.2 million of these were not limited and had "ifconfig" on them so they could be identified
  - 420 Thousand of these met minimum CPU/RAM requirements of the researcher and were therefore used to perform Internet Census 2012
  - A list of these devices has never been published

AusCERT

What is the Internet Census 2012?

- Complete Scan of the allocated IPv4 ranges of the Internet
- Results were released Mid-March by an anonymous researcher along with a paper
- http://internetcensus2012.bitbucket.org/paper.html
- Results contain 9 TB of logs (pure text!)
- Publicly available for download through a torrent as 568 GB of highly compressed (ZPAQ) files
- Details on my thesis project on the Internet Census: http://bit.ly/census-project-proposal

What is in the 9 TB of data?

- ICMP Ping (52 billion records) - 1.8 TB
- Reverse DNS (10.5 billion records) - 366 GB
- Service Probes (175 billion records; 4000 billion requests) - 5.5 TB
- Hostprobes (19.5 billion records) - 771 GB
- Syncscans (71 billion ports scanned) - 435 GB
- TCP IP Fingerprint (80 million records) - 50 GB
- IP ID Sequence (75 million records) - 2.7 GB
- Traceroutes (68 million records) - 18 GB

AusCERT

How is this feasible?

- Maximum of 4,294,967,296 IPv4 Addresses
  - Only 3,706,650,624 are allocated
- Using only 1 device to scan and performing comprehensive a scan of 1 IP per second, it would take:
  - 3.7 billion seconds ≈ 117.5 Years
- But with 420,000 devices it would only take 2.6 hours!
- In under 24 hours you can easily collect all the data you need for all allocated IPv4 ranges!
- For problems of logistics and how the researcher handled collection of the data refer to the Internet Census 2012 paper

AusCERT

How to be part of the Carna Botnet?

- A device must be directly reachable from the Internet
- Telnet running on default port 23 (with no firewall for protection)
- Allow login using one of the default credentials
  - E.g. admin:admin, admin:password, root:password etc

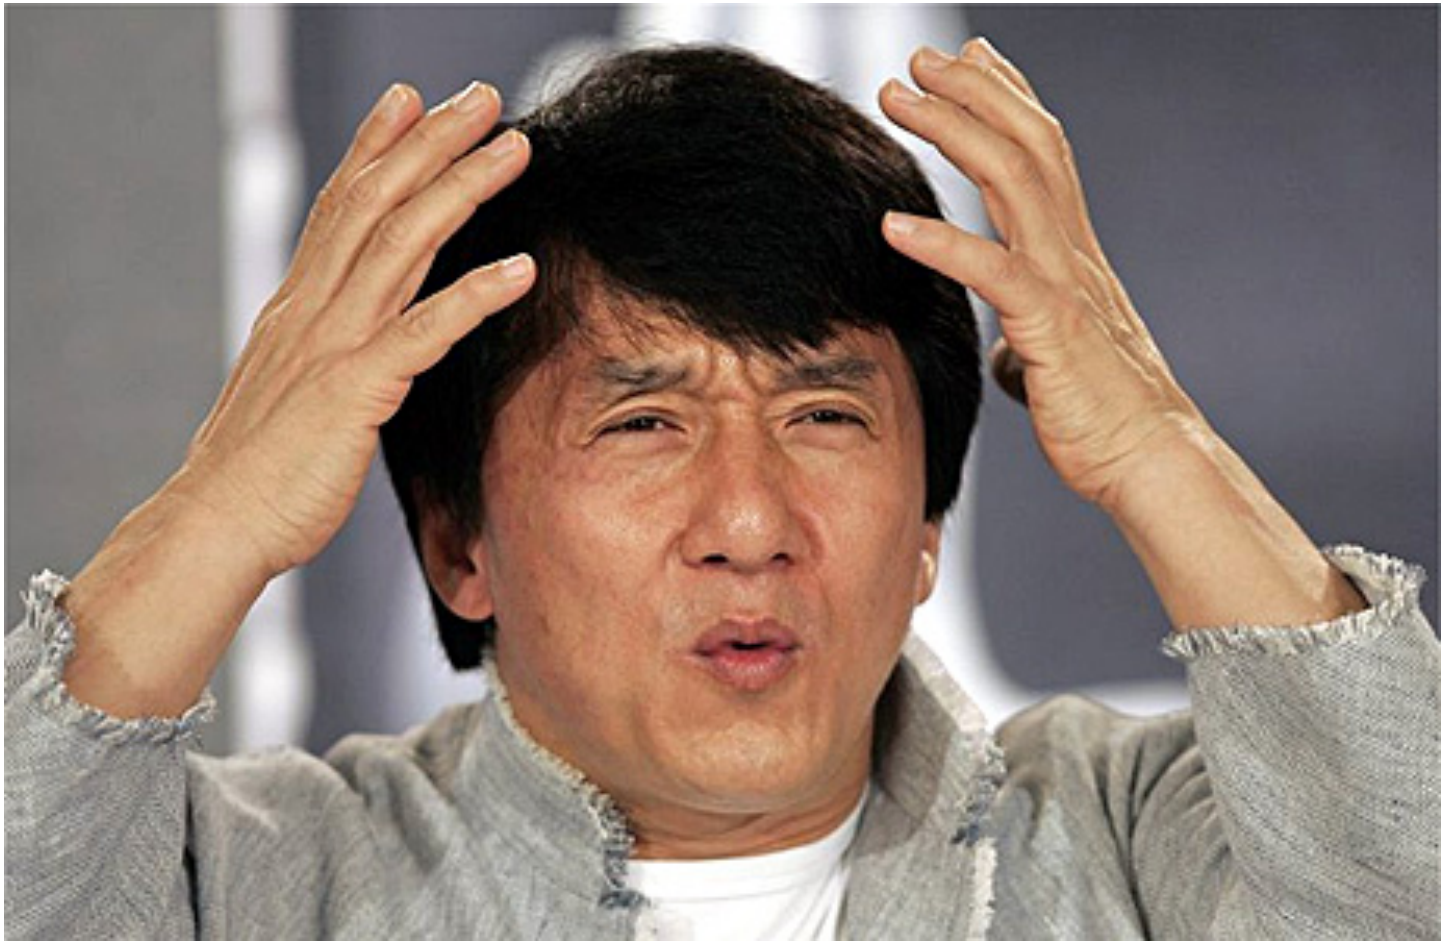

by Parth Shukla on 2013-08-27 @ APNIC 36 - Xian, China

How to be part of the Carna Botnet?

- A device must be directly reachable from the Internet
- Telnet running on default port 23 (with no firewall for protection)
- Allow login using one of the default credentials
  - E.g. admin:admin, admin:password, root:password etc
- Not just make 1 mistake but 3 mistakes to be part of this botnet!
- To be part of the 1.2 million analysed here, also needed 'ifconfig'
- To be part of the subnet of 420k further needed ability to upload custom binary and have met some minimum RAM and CPU specs so as to avoid interfering with industrial controls or mission critical hardware

AusCERT

This presentation

- About the ≈1.2 million identifiable compromised devices
  - This data obtained directly from the anonymous researcher
  - Used for analysis for the rest of the presentation
  - NOT publicly available!
- From now on Carna Botnet = 1.2 million identified devices
- Particular focus on devices located in China in the data
- This botnet is unusual because it's not created by phishing, exploiting a coding error or social engineering!

Why so many devices?

- ≈1.2 million! WHY?!
- Are there really that many 'stupid' people?
- We will come back to this later

- Let's develop some foundation and context first

AusCERT

Information for each device

1. **MAC address** - the last byte replaced by an ascending number
2. **Manufacturer** - derived from MAC address
3. **RAM** - in kilobytes as that's in /proc/meminfo
4. **Uname** - output of uname -a
5. **CPU Info** - output of /proc/cpuinfo
6. **IPs** - list of all IPs associated with this device. Last byte of each IP was zeroed by researcher. Accuracy of each IP to within a C class.
   - **Country Code** - two letter country code for each of the IPs. Correct at the time the device was compromised.

AusCERT

In the 1,285,192 devices there were

- 200 unique country codes
- 2,098 unique device manufacturers
- 3,880 different RAM sizes
- 10,875 unique unames
- 35,997 unique CPUs
- 787,665 unique C class IP ranges
- 1,264,223 unique MAC addresses

AusCERT

Device Distribution by Countries

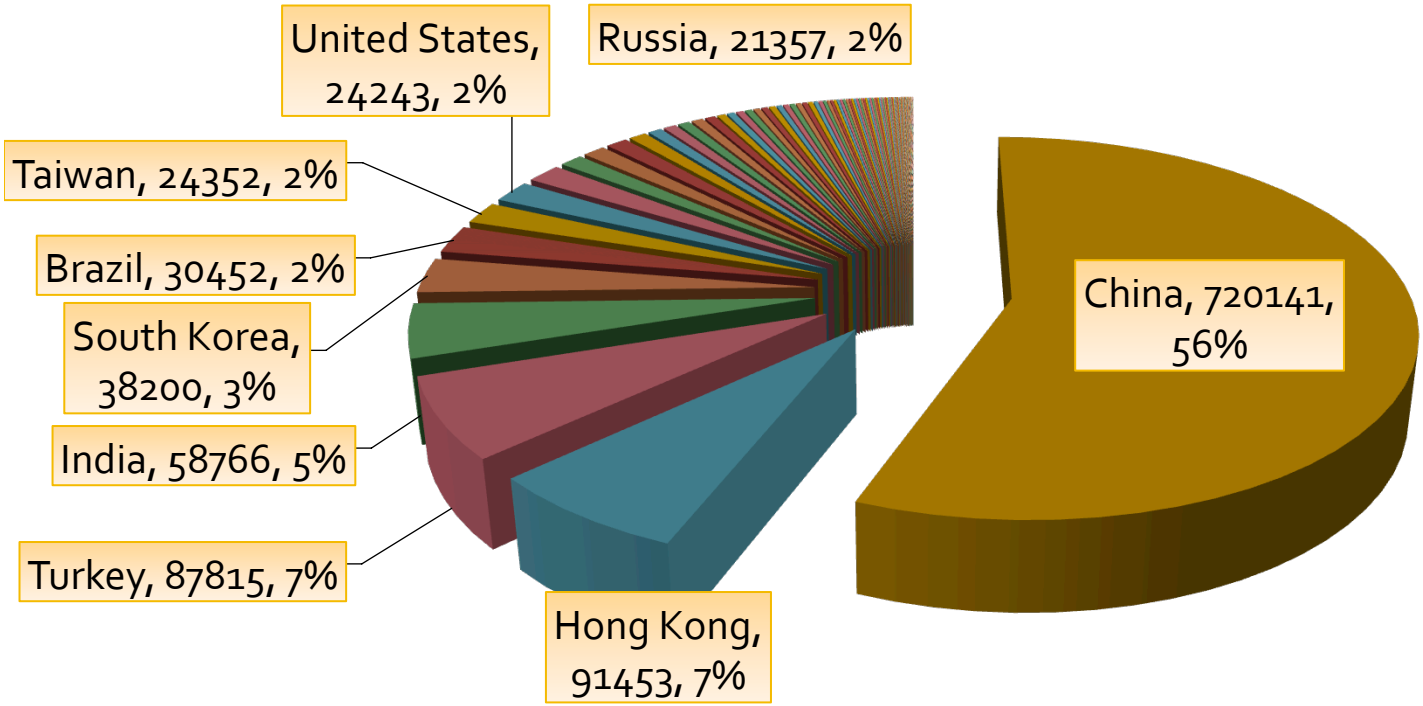

AusCERT

Device Distribution by Continents

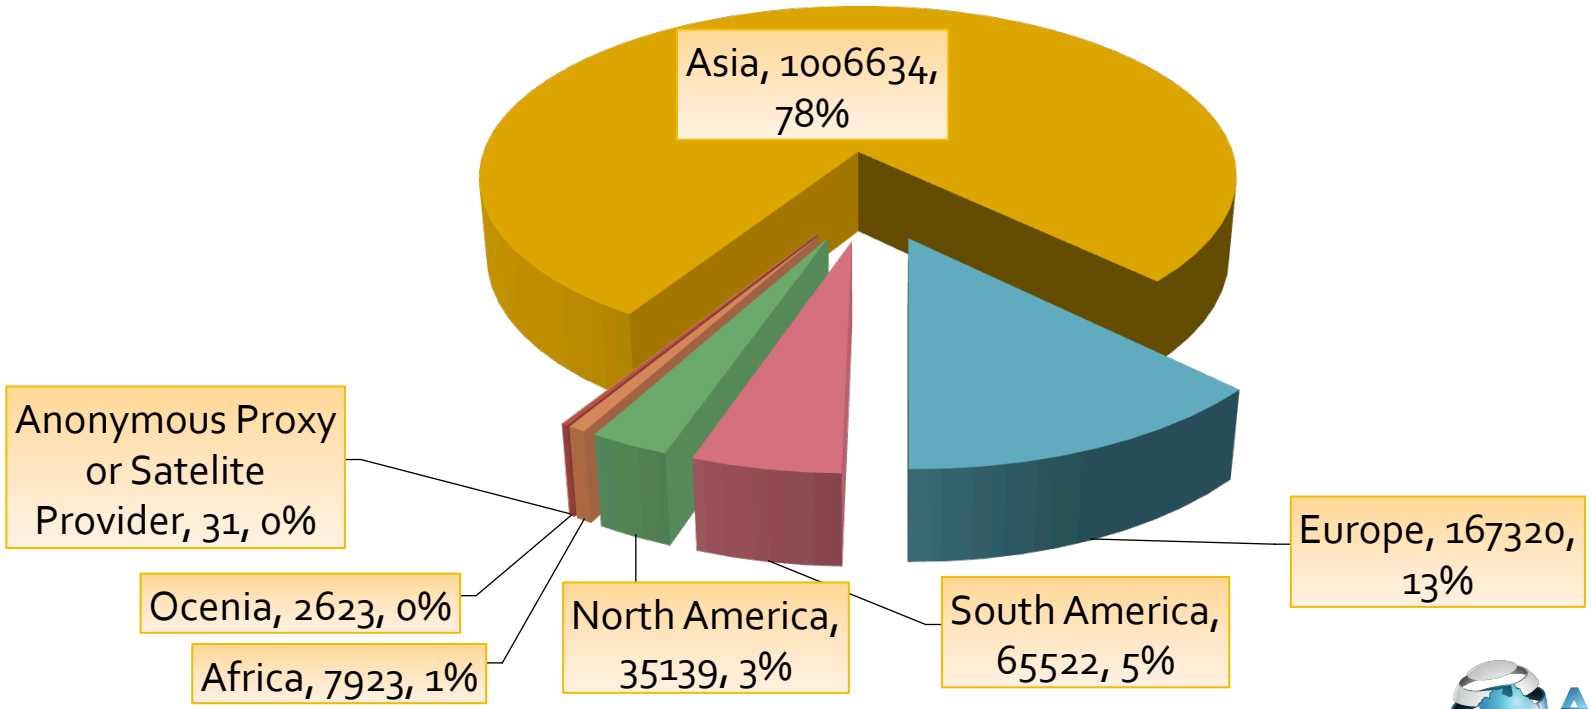

AusCERT

Device Distribution for Asia

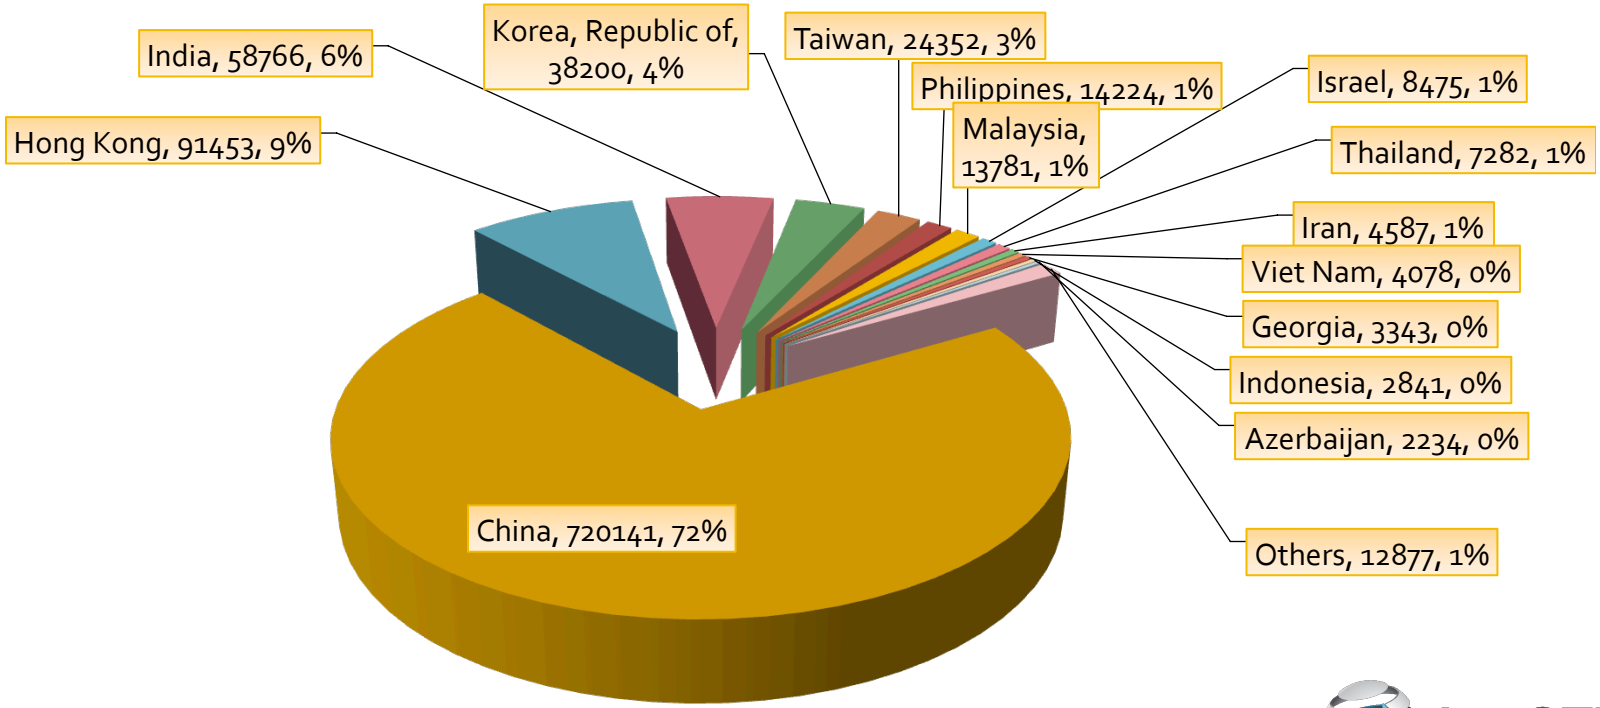

AusCERT

Where is China?

- 720,141 devices located in China in the data
- Largest Slice (56%) Worldwide
- Also Largest Slice (72%) in Asia

- Numbers should be terrifying given the prevalence of vulnerable/ infected devices in China!
- For a worse scare, we will look at how easy it would be to find one of these devices in China at the end of the presentation.
  - China is not the worst affected despite the large numbers

Distribution of Manufacturers - Worldwide

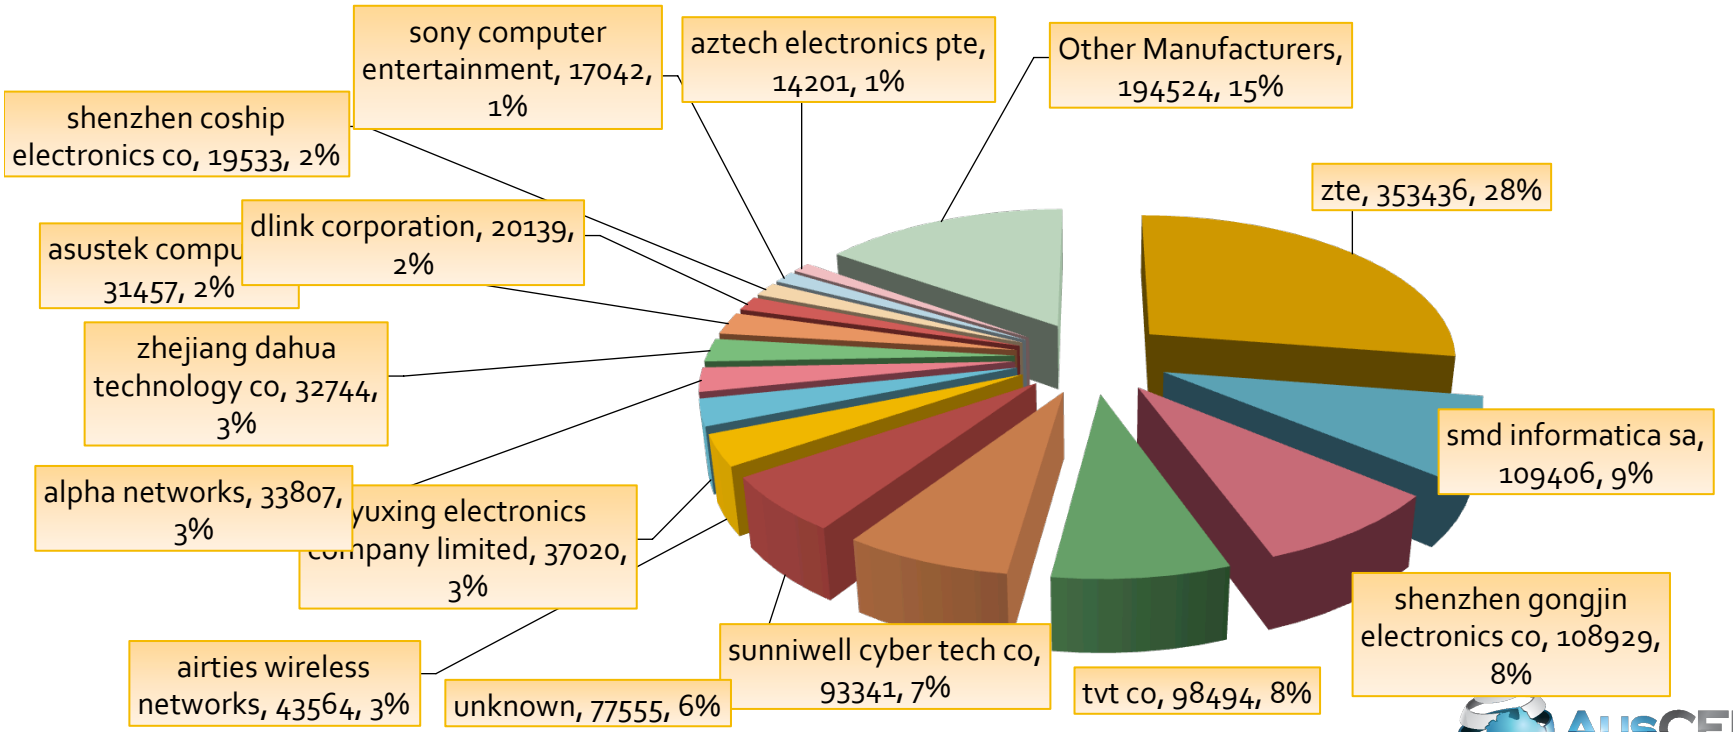

AusCERT

Distribution of Manufacturers - Asia

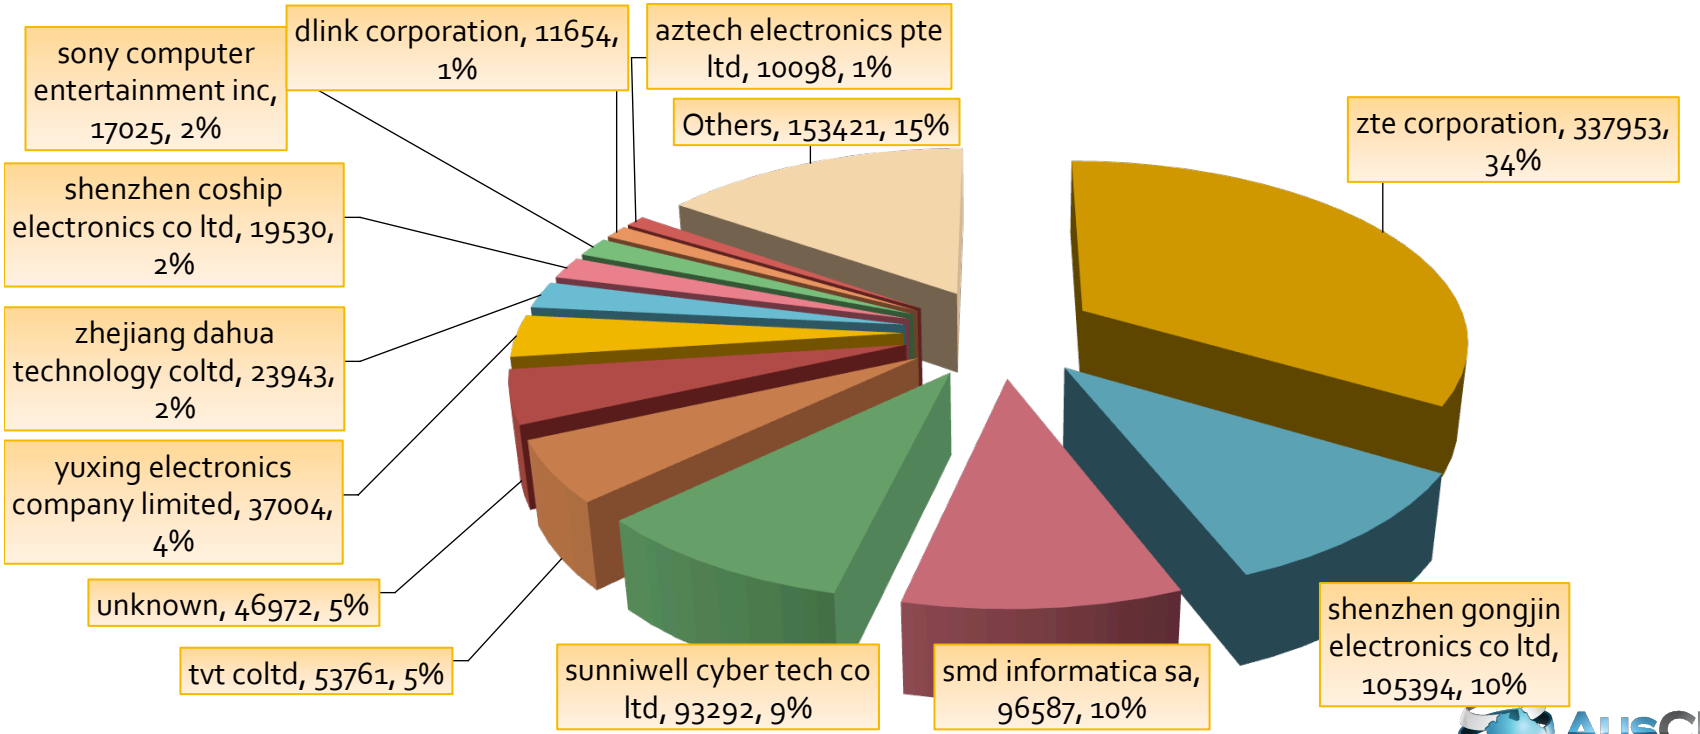

AusCERT

Distribution of Manufacturers - China

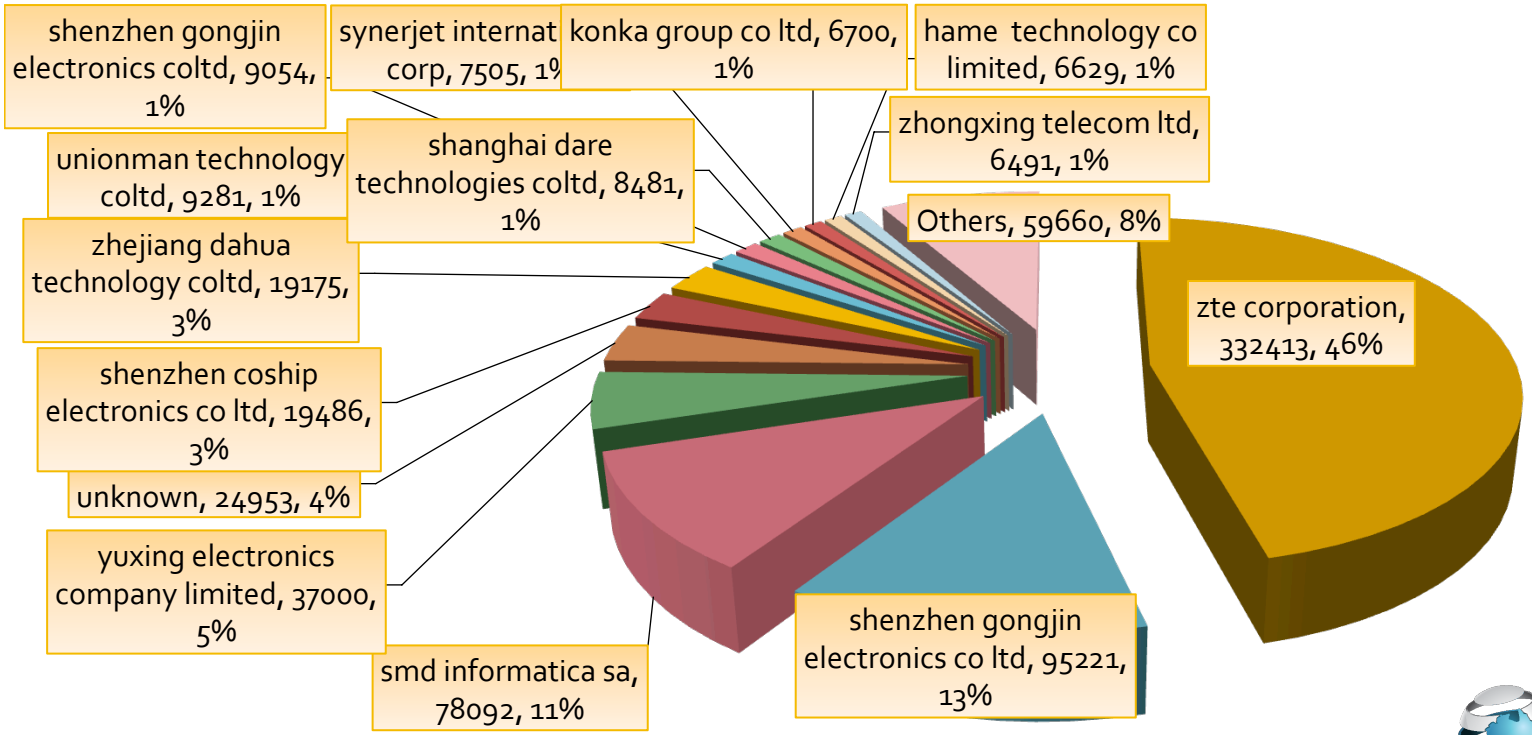

AusCERT

What do these companies sell?

- ZTE – Chinese Company
  - Mobile Phones
  - Hardware, software and services to telecommunications providers
- Shenzhen gongjin electronics – Chinese Company
  - Home Networking Products (Modems & Routers)
- SMD INFORMATICA S.A. – Portuguese Company
  - No information in English that I could locate!

AusCERT

Number of Manufacturers

- Worldwide - 2,098 unique device manufacturers
  - Can only see 14 in the Graph – rest in the "Others" category
- Asia – 1,643 unique device manufacturers
  - Can only see 11 in the Graph – rest in the "Others" category
- China – 1,111 unique device manufacturers
  - Can only see 13 in the Graph – rest in the "Others" category
- Above does not count the "unknown" category
- Why are the rest of the manufacturers so small?
- More importantly why are some SO BIG!

AusCERT

Number of Manufacturers by Number of Records

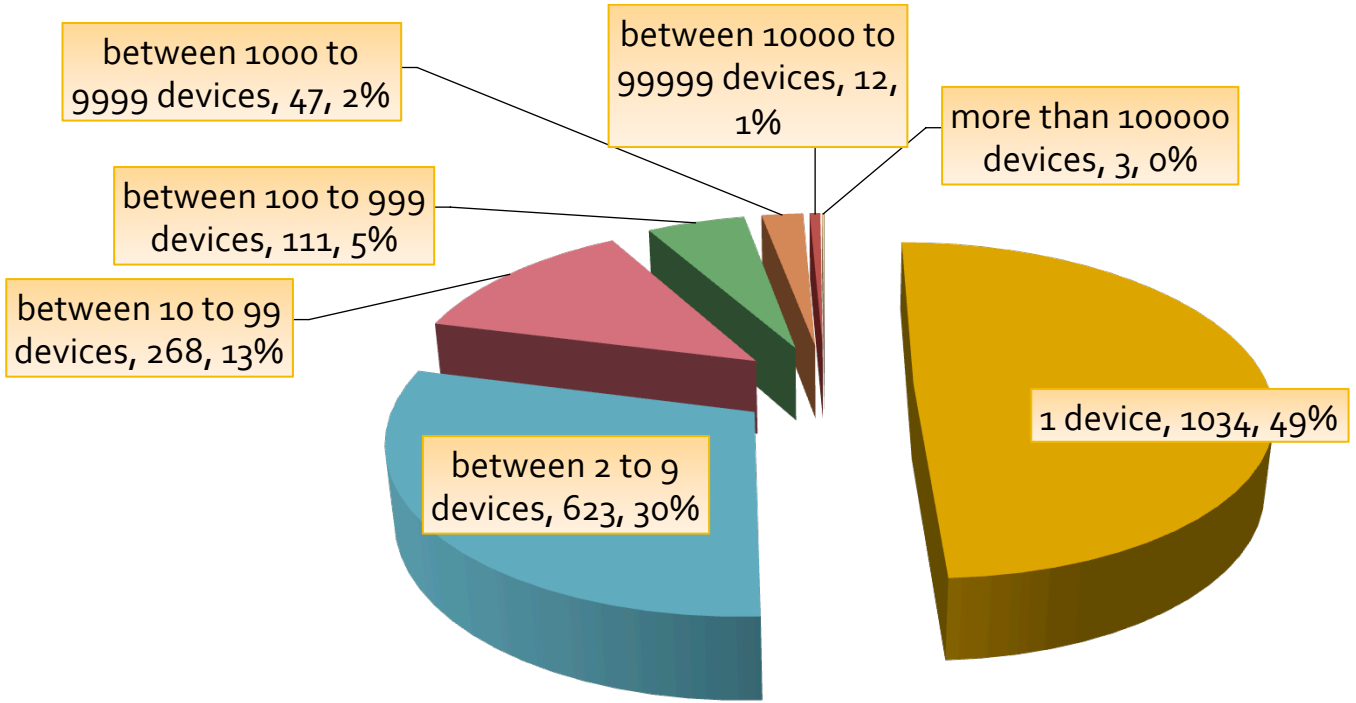

AusCERT

Asia - Number of Manufacturers by Number of Records

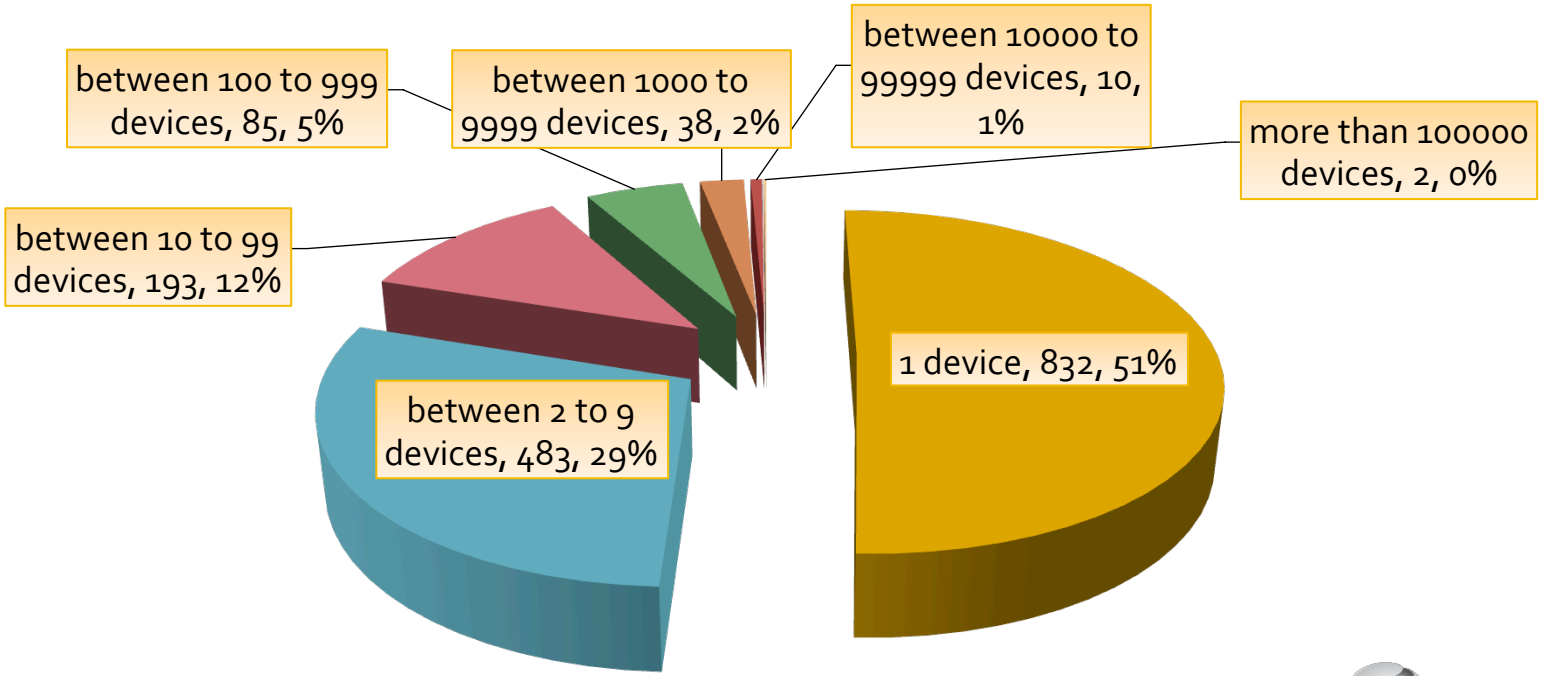

AusCERT

China - Number of Manufacturers by Number of Records

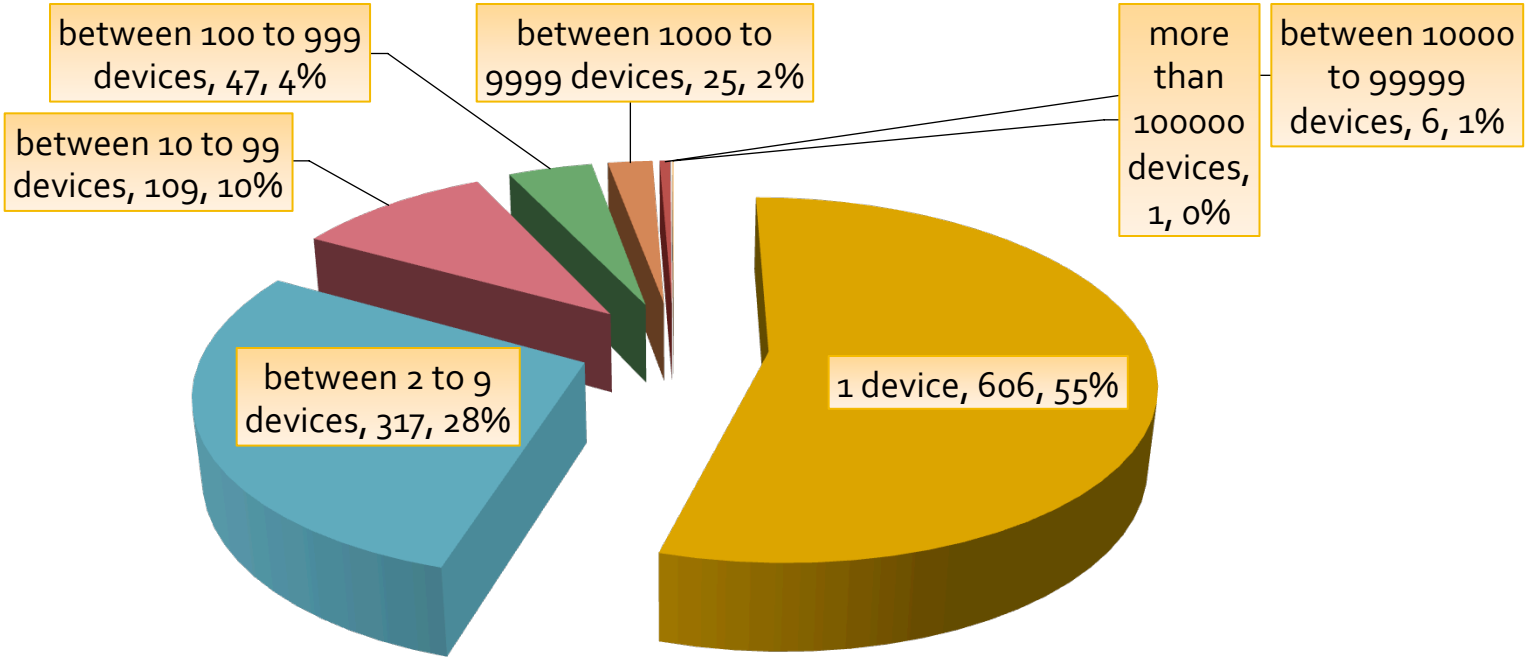

AusCERT

So why so many devices?

- Given the prominence of certain manufacturers, it seems obvious that most devices in the data are not because of 'stupid' people
- Certain devices by certain manufacturers may:
  - not allow the change of default logins for telnet
  - Have a 'backdoor' hardcoded with default credentials perhaps to allow for remote diagnostics (ISPs could have requested this!)
  - Lack of documentation that there is even a telnet server running on it!
    - what device wouldn't you bother looking for an open telnet port?
  - Require devices to have Internet Reachable IP to benefit from full functionality of the product (i.e. remote viewing of CCTV camera)

Major Manufacturers from where?

- Global Devices:
  - 6 out of 14 visible manufacturers based in China
- Asian Devices
  - 6 out of 11 visible manufactures based in China
- Chinese Devices:
  - Most of visible manufacturers based within China

AusCERT

Progress of presentation

- Data fields analysed so far:
  - Countries
  - Manufacturers (& MAC Addresses)
- Still to have a look:
  - Uname – Very quick look next
  - RAM – Up next
  - CPU Info – Not analysed due to inconsistency of the field
  - IP Addresses – at the end for a scary ending

AusCERT

Worldwide Distribution of 'uname'

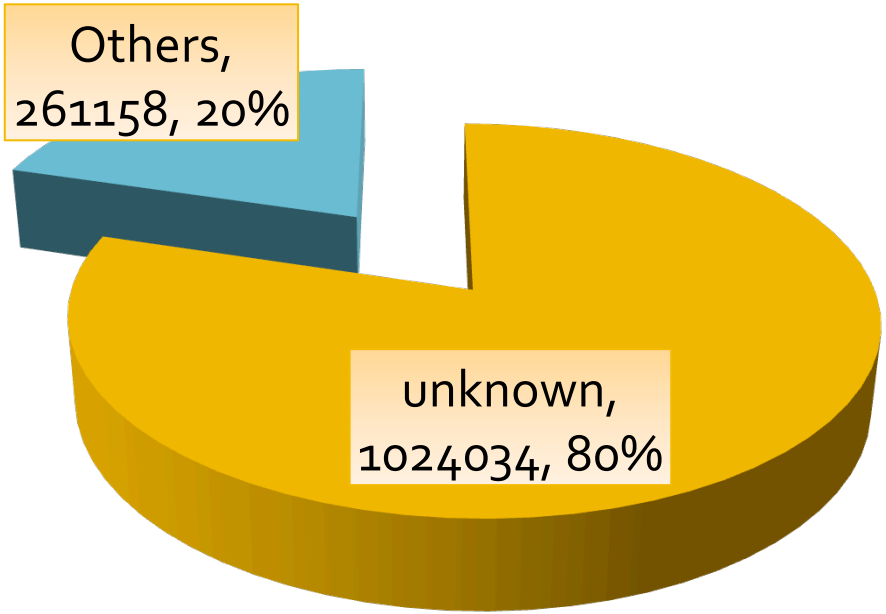

AusCERT

Worldwide Devices – RAM Distribution

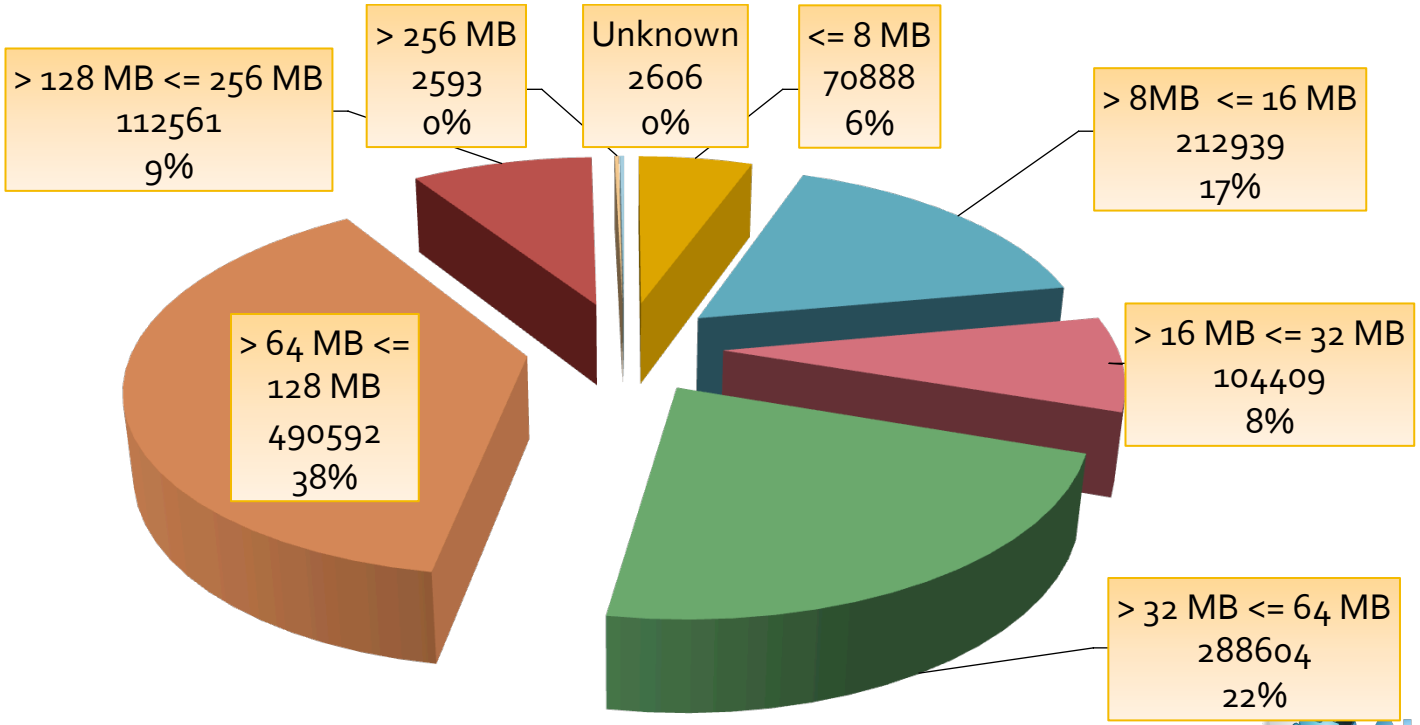

AUSCERT

Worldwide Devices - RAM Statistics

Description	Value
Unique RAMs	3,880 different RAM sizes
Lowest RAM	5,488 kilobytes (5.35 MB) – 1 device in Germany
2nd Lowest RAM	5,688 kilobytes (5.55 MB) – 1 device in USA
Highest RAM	4,828,263,435 kilobytes (4.49 TB) – 1 device in China
2nd Highest RAM	1,000,000,000 kilobytes (0.93 TB) – 5 in China, 1 in Ukraine
Most common	11,500 kilobytes (11.2 MB) – 98,947 devices (7.7%)
2nd Most common	124,620 kilobytes (121.7 MB) – 96,543 of devices (7.5%)

AusCERT

Asian Devices – RAM Distribution

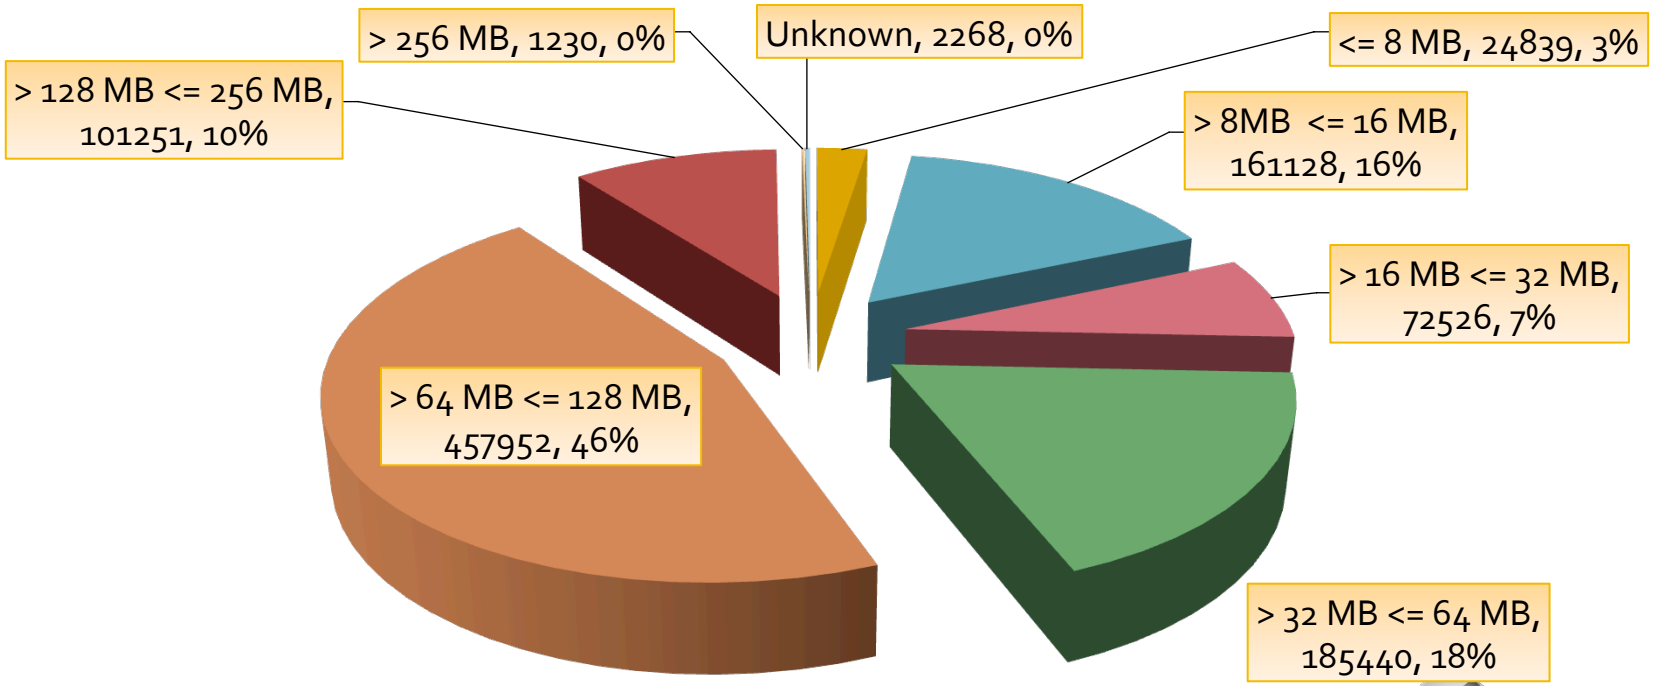

AusCERT

Asian Devices – RAM Statistics

Description	Value
Unique RAMs	2,612 different RAM sizes
Lowest RAM	5,700 kilobytes (5.56 MB) – 1 device in China
2nd Lowest RAM	5,752 kilobytes (5.61 MB) – 1 device in Taiwan
Highest RAM	4,828,263,435 kilobytes (4.49 TB) – 1 device in China
2nd Highest RAM	1,000,000,000 kilobytes (0.93 TB) – 5 devices in China
Most common	11,500 kilobytes (11.2 MB) – 98,875 devices (9.8%)
2nd Most common	124,620 kilobytes (121.7 MB) – 96,543 of devices (9.6%)

AusCERT

Chinese Devices – RAM Distribution

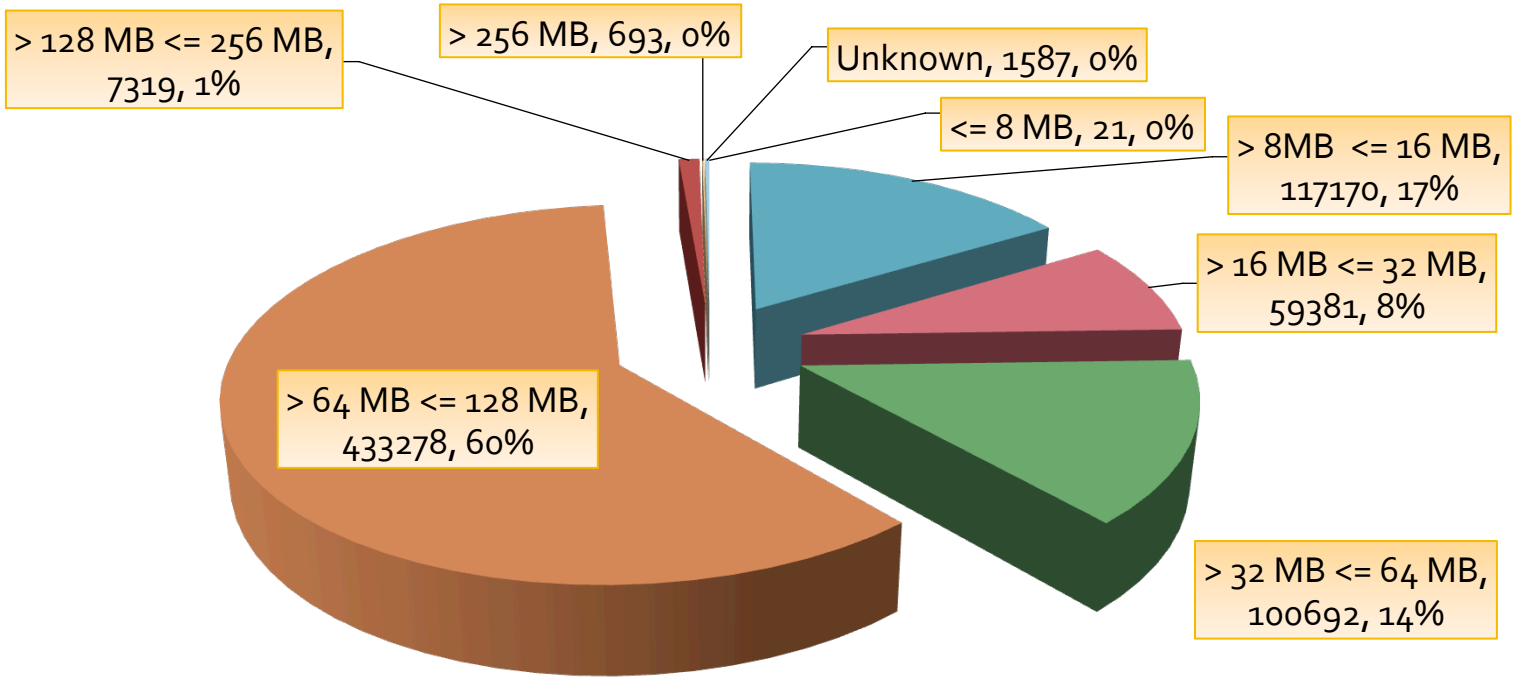

AusCERT

Chinese Devices – RAM Statistics

Description	Value
Unique RAMs	1,551 different RAM sizes
Lowest RAM	5,700 kilobytes (5.57 MB) – 1 device
2nd Lowest RAM	6,128 kilobytes (5.98 MB) – 1 device
Highest RAM	4,828,263,435 kilobytes (4.49 TB) – 1 device
2nd Highest RAM	1,000,000,000 kilobytes (0.93 TB) – 5 devices
Most common	11,500 kilobytes (11.23 MB) – 98,743 devices (13.7%)
2nd Most common	124,620 kilobytes (121.7 MB) – 96,543 of devices (13.4%)

AusCERT

Worldwide RAM vs Asian RAM

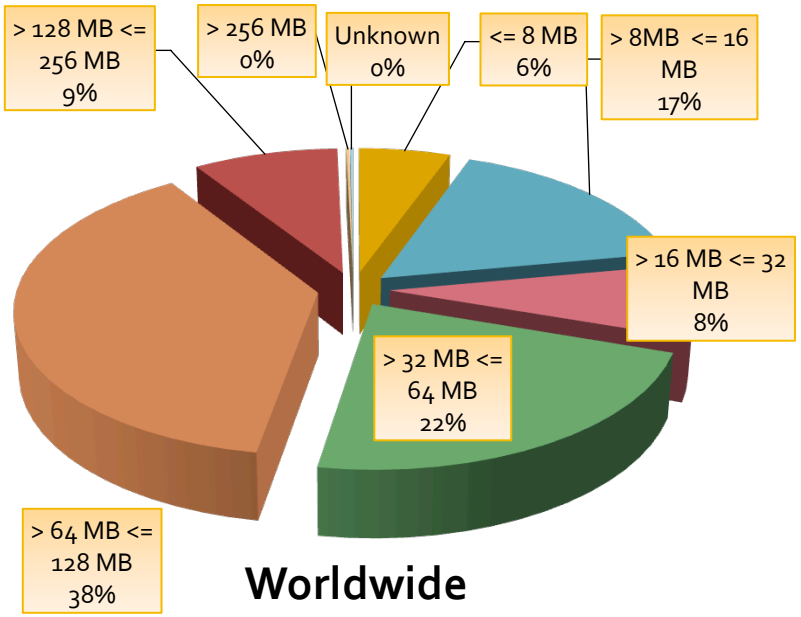

Worldwide

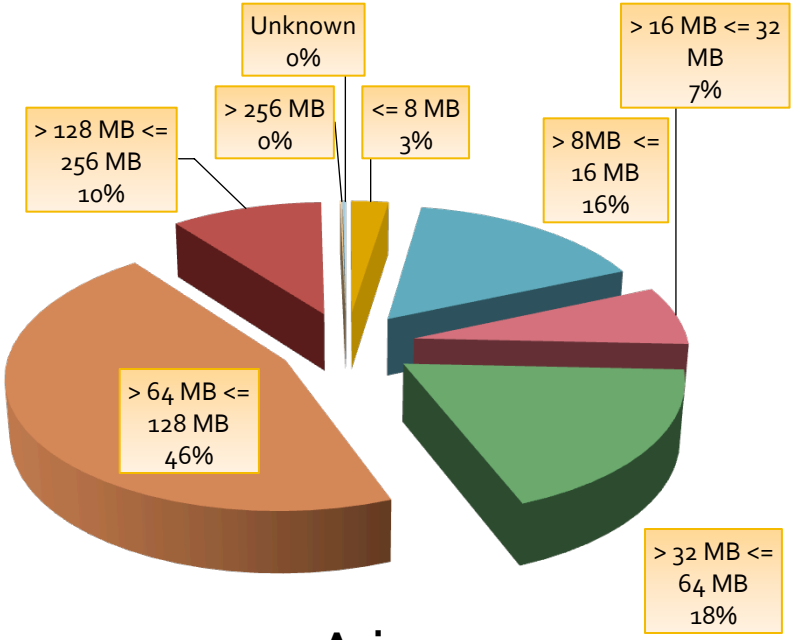

Asia

AusCERT

Asian RAM vs Chinese RAM

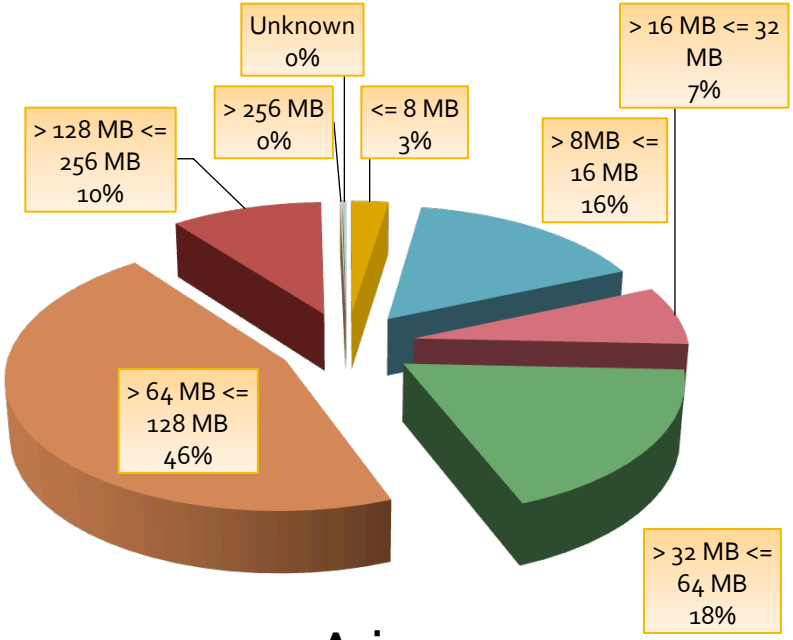

Asia

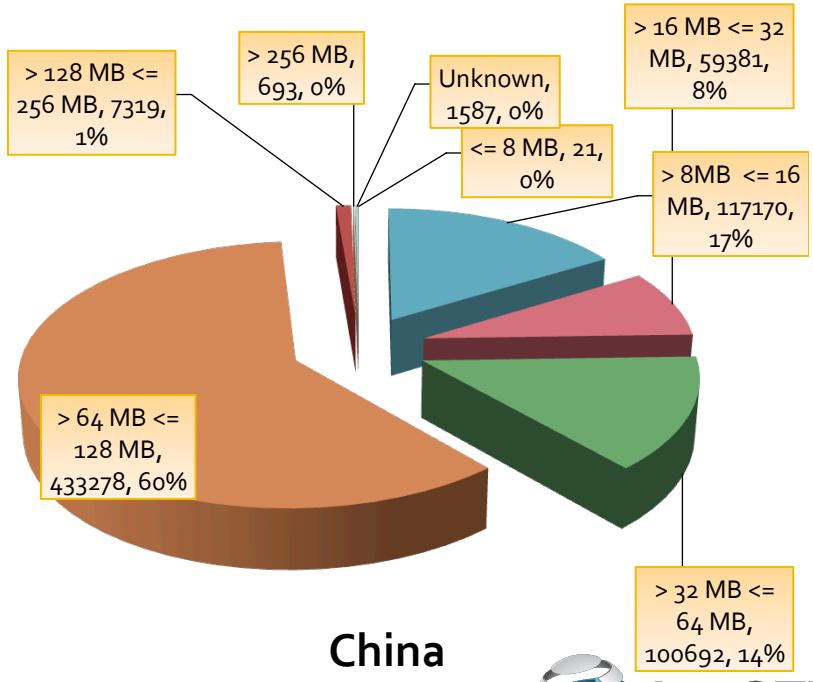

China

AUSCERT

Worldwide RAM vs Chinese RAM

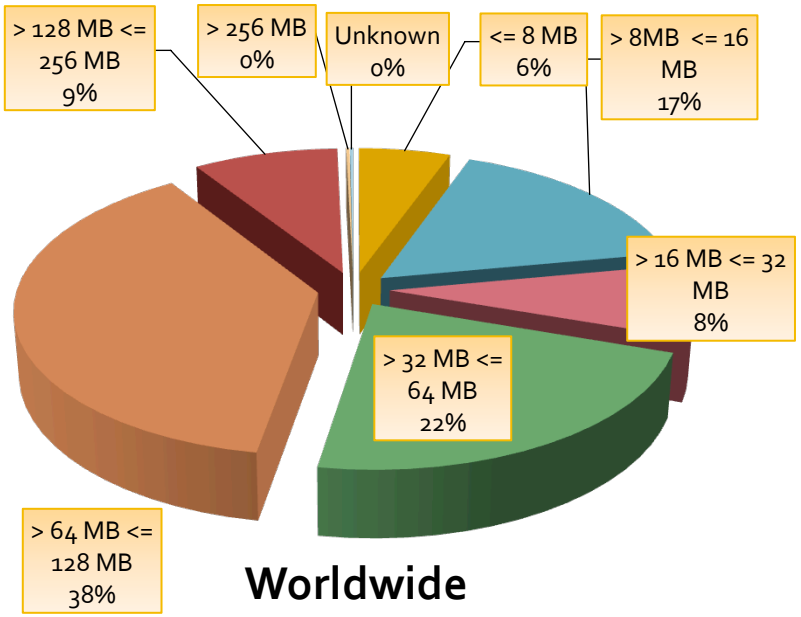

Worldwide

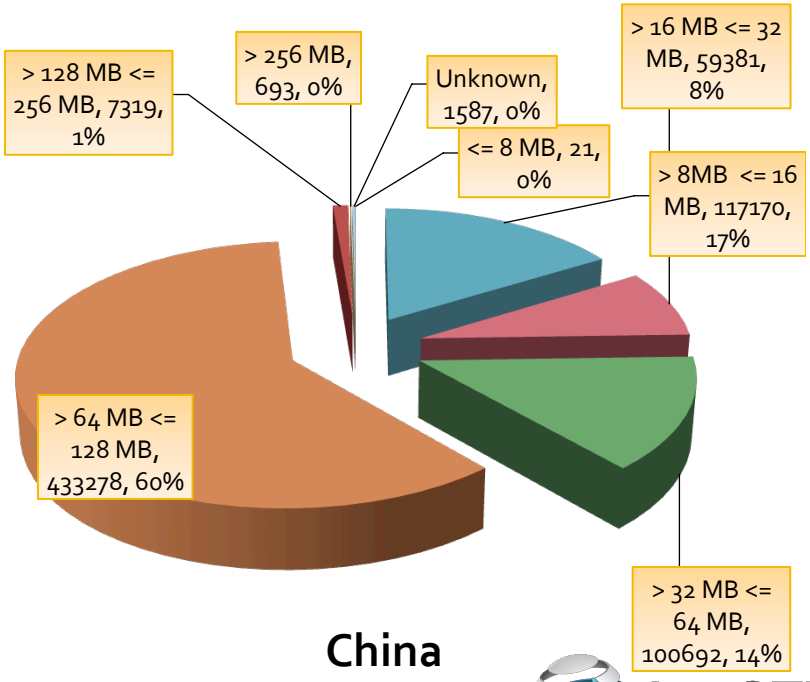

China

AUSCERT

Scary Stats – How easy to find a device?

- We can calculate how easy it would be for someone to find a vulnerable device with this simple equation:

$$= No.\ of\ infected\ devices\ for\ the\ region / No.\ of\ Allocated\ IP\ ranges\ for\ the\ region$$

$$= Infected\ devices\ per\ IP\ range\ for\ the\ region$$

AUSCERT

Clarification on Calculations

- IP Allocations have changed over time
  - World/Asia or China did not have as many IPs allocated during the initial formation of the Carna Botnet as it does now
- Old IP allocation statistics from RIRs were used in these calculations
  - Allocated IP ranges as of 1st December 2012 were used for calculations to get an accurate idea of infection ratio around the time the Carna Botnet was formed
- Rations are assumed to be a good approximation for now as well
  - assuming that the rate of growth of allocated IPs is directly comparable to rate of growth of vulnerable devices added to the Internet over time

Worldwide - How easy to find a device?

No. of infected devices for the World/No. of Allocated /24 IP ranges for the World = 1,285,192/13,587,587

- ~0.095 device per /24 IP range
- ~9.46 devices per 100 C class ranges
- Average 1 vulnerable device every ~10.57 subnet
- Average 1 vulnerable device every ~2706 IPs
- Scanning 10 IPs/sec would take ~4 minutes 31 seconds to find a device
- No. of Allocated /24 IP ranges for the world deduced by adding all allocated ranges by each of the Regional Registries as of 1 December 2012

AusCERT

Asia - How easy to find a device?

No. of infected devices for Asia/No. of Allocated /24 IP ranges for Asia = 1,006,634/3,260,028

- ~0.309 device per /24 IP range
- ~3.09 devices per 10 C class ranges
- Average 1 vulnerable device every ~3.24 subnets
- Average 1 vulnerable device every ~829 IPs
- Scanning 10 IPs/sec would take ~1 minute 23 seconds to find a device
- No. of Allocated /24 IP ranges for Asia deduced by adding all allocated IP ranges for each country in Asia as of 1 December 2012

China - How easy to find a device?

No. of infected devices for China/No. of Allocated /24 IP ranges for China = 720,141/1,289,054

- ~0.5587 device per /24 IP range
- ~5.59 devices per 10 C class ranges
- Average 1 vulnerable device every ~1.79 subnets
- Average 1 vulnerable device every ~458 IPs
- Scanning 10 IPs/sec would take ~45 seconds to find a device
- No. of Allocated /24 IP ranges for China deduced by adding all allocated IP ranges for China from APNIC as of 1 December 2012.

AusCERT

Terrified now? But wait!

- So just probe these devices on port 23 to find them?
- NO, because a scan of some of the IP range from the data show almost all ranges had port 23 closed?! Faulty Data?
- Carna Botnet shutdown telnet to close port 23 as soon as it had control of the device and setup iptable rules where possible
  - Primarily to avoid interference from other botnets
  - Temporary only - settings lost on reboot
  - Other botnets won't be so nice
- If telnet is the only shell into the device then hardware reset is the best chance of ensuring a clean device

It gets worse!

- 1308 IP ranges were found that appear in more than or equal to 260 different records.
  - Same IP range in more than 260 different device records
  - i.e. almost all IPs within these 1308 IP ranges likely to contain vulnerable devices!
  - This is not a 'guarantee' as devices compromised at different times
- If you were to find these 1308 IP ranges and "hog" them then you'd have a botnet of: ~327 thousand
- Which countries more prominent in these records?

AusCERT

Countries of IP ranges appearing in more than 259 records

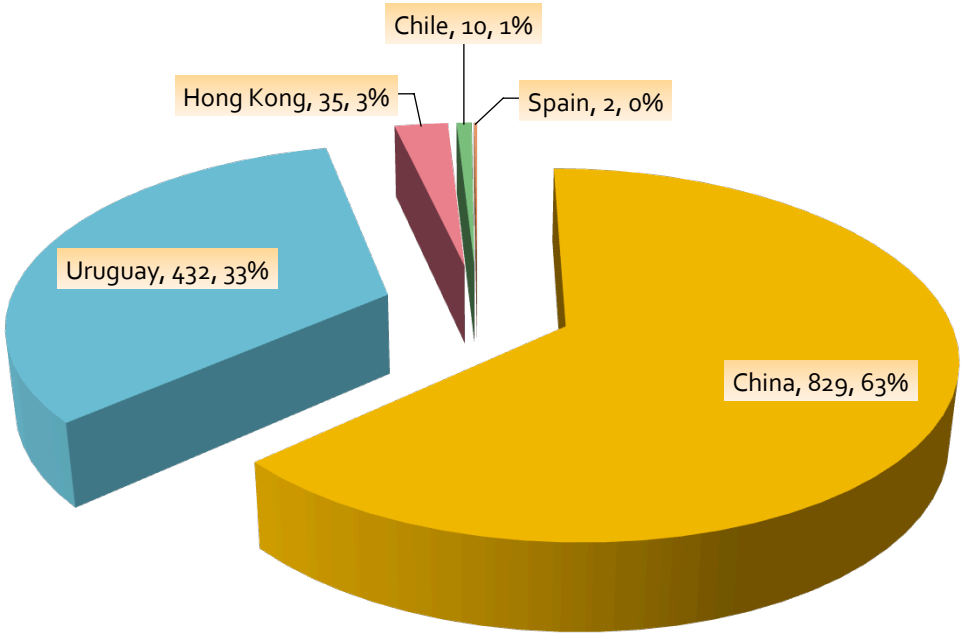

AusCERT

Scariest News

- Open source tool 'lightaidra' does exactly what Carna botnet did:
  - Auto searches for telnet ports with default creds
  - Allows you to upload your custom binary that can do anything!
    - For routers it can sniff traffic, modify traffic! Spam the world! Anything!
  - Joins IRC chat room to read latest commands
- Bad guys really don't need to do much
- Carna detected presence of Aidra (as noted in the paper)
  - So this data might not be ALL vulnerable devices
  - However, Carna was a lot more cross-platform than lightaidra is by default
- Info: http://vierko.org/tech/lightaidra-0x2012/

A Glimmer of Hope

- Most embedded devices mount their partitions as read-only. tmp and other directories stored in RAM
- For most devices a reboot would lose the malware
  - May leave port 23 closed!
  - Start up scripts (if functionality exists) could've been modified to re-infect
- Almost all the time a hardware reset and/or firmware reflash will resolve the problem
  - If malware authors wanted they could interfere with re-flashing and hardware reset depending on how much control telnet allowed over the device

Any Good News for China?

- China does not have the worst infection ratio for a coutry/region in the World. There are other countries that have worse!
  - Hong Kong has worse ratio then China!
- If you want to find out which countries are worse please read my research paper
- China CERT have the Chinese data and are working to make a difference for China!

AusCERT

What have I done?

- Supplied relevant data to APCERT members and to CERTs from any country with more than 10 thousand compromised devices
  - Including CNCERT/CC
  - Adds up to about 26 countries in total
- Split Australian data by ISP and provided to them so they can deal with the problem on their own network
- Contacted IEEE to contact the worst affected manufacturers
  - Reached out to manufacturers to work with us; only 1 of 23 has responded
- Presentations and Research Paper!
  - Talking to people and publication of presentations and research paper

Research Paper

- Data wasn't provided on a silver platter ready for analysis
- Checking for consistency both internally and externally was done
- Manufacturer field was re-derived
- Assumptions had to be made; duplicates were removed
- Efforts to check for accuracy were made
- Contains a complete list of countries, infection ratios and manufacturers
- All of this plus more is covered in detail in the Research paper at: http://bit.ly/carna-paper

Problem needs attention NOW

- Devices behind NAT not in the data
  - Researcher did not scan internal network when a router was compromised so.
  - So number of actually vulnerable devices likely to be a lot bigger
- With IPv4 to IPv6 transition happening now, this is the time to make sure such devices are secure by default.
  - Specially since NAT "protection" will not be available on IPv6 and
  - bad router firewalls might expose even more devices then visible in this data set.

AusCERT

Public Awareness

- Seemingly harmless/small threat is real and big!
- Awareness of problem is the first step
- Participation of diverse range of players from the industry required
- Awareness to be raised with public/manufacturers/ISPs-selling-the-devices on the problem of port 23 open with default login
- This presentation is one of the many steps required to tackle the problem as a whole
- Please spread the word!

I can't do it alone. What can you do?

- It's a long and hard battle because no 'easy' or 'quick' solutions
- Read my detailed Research Paper
- Re-look at these slides online
- Tell people: family, relatives, employer, friends, colleagues
- Secure your devices! Secure others' devices
- Do you know anyone in a position who can help? Maybe you are?
- Help influence government, companies, ISPs, manufacturers into ensuring devices used and sold by them are secure by default
- Have ideas on tackling the problems? Contact me!

Research Paper: http://bit.ly/carna-paper

This presentation: http://bit.ly/carna-apnic