RESTing on our laurels

AUTOMATING THE NETWORK THE "EASY" WAY

Ayden Beeson

Senior Technology Specialist (Networks)

Charles Sturt University

It all began with a dream

The “hammocks for all” network

- Standardised devices, including port count
- No ”special” switch configurations (”sheep not pets”)
- Automated switch provisioning – no more copying configs
- Automated ports – no more manual VLAN changes

Standarised models

- Cisco 3850's selected for building access switches
- 48 ports only
- POE only
- Mgig ports available
- 10 Gig SM Optics only (with 1 exception)
- 3850-12X48U met all requirements

Standarised models

- Cisco 3560s selected for small location access
- 8 ports only
- POE only
- Mgig ports available
- 10 Gig SM Optics where possible, Mgig otherwise
- 3560CX-PD met all requirements

Standarised models

- Cisco 6800 series selected for Core/Distribution
- Fixed chassis and modular options
- VSS supported
- SFP+
- 6840's used where possible, 6807's where density required
- 6880's used in the DC for density and XL BGP tables
- Unfortunately IOS firmware files not common across models

Standarised models

- Cisco 43xx ISRs selected for remote sites w/ VPN links back
- High throughput crypto
- Additional throughput license if required
- Additional switching line card supported for very small sites
- Modules available for VDSL, 4G, etc as required
- 4331 / 4321 met the requirements
- 4331 for rack installs, 4321 for desk/table installs

OK so we have standard models,
but now what do we do with them?

No ”special” switch configurations

- Original network had:
  - 38 Management vlans
  - 27 static vlans
  - 47 Staff vlans
  - 25 lab vlans
  - Lots of various other entries as required

No ”special” switch configurations

- Consolidated L3 routers down to single core at each site
- Short DHCP lease timers to smooth transition timers
- Moved hosts and IPs as required
- Condensed multiple /24 Vlans into single larger pools
- All done using old 3750’s and 6513 core switches

No ”special” switch configurations

- New network has:
  - 1 Management vlan
  - 1 static vlan
  - 1 staff vlan
  - 1 lab vlan
  - 1 of every other type we required

No ”special” switch configurations

- All new switches to connect straight back to the core
- No active fibre daisy chaining
- Required an investment to remove legacy MM fibre where required
- 1 large site at Wagga Wagga required 2 active fibre distribution 6800s for geographic reasons

Now we have a standard configuration to put on every switch, but we need to get it out there….

Automated switch provisioning

- Prime configuration templating set up for 3850 and 3560s
- Allowed us to tailor the configuration per site / per stack
  - Device name
  - Management IP (IPv4 and IPv6)
  - SNMP Location
  - Stack member counts + port configurations as required

Automated switch provisioning

- Prime Infrastructure integrates into APIC-EM
- Prime sets up the APIC-EM Projects and deploys matching firmware
- Prime API PnP support very limited at the time (PI 3.0)
- Added in PI 3.1 to allow pushing devices in directly
- Will get added into our workflow very soon

Automated switch provisioning

- Switches provisioned via APIC-EM's PnP module
- APIC-EM handles:
  - Switch "onboarding" from a provisioning VLAN
  - Pre-provision of devices from Prime with known serials
  - Software updates to "expected" version
  - Crypto key generation
  - Configuration deployment (from Prime templates)
  - Verification of all of the above and retries if not successful

Automated switch provisioning

- Once the switch is imported in Prime
  - A configuration is generated according to the template, using the variables provided
  - This configuration is pushed to APIC-EM with the serial number of the switch
  - The switch itself is then inserted into APIC-EM as a pre-provisioned entry
  - Whole sites can be pre-deployed this way, ready for cutover

Automated switch provisioning

- Switch provisioning can be done in the field, out of the box
- Default trunks must have provisioning vlan as native (usually vlan 1)
- Switches get IP and APIC-EM settings via DHCP option
- Boot up, upgrade and provision of configuration takes around 30-60 minutes all up

Great, all my switches are the same,
now how can I keep it that way?
Where is the network *magic?*

Automated ports – The magic

- Using Cisco's IBNS (Identity-based networking services) to handle advanced use cases
- Using 802.1x and MAB to provide dynamic VLAN assignment
- Cisco ISE is used to parse the RADIUS sessions and provide the required VLAN and SGT options back to the switch
- Port policy stored in templates to keep interfaces "clean"

Automated ports – The magic

- IBNS gives us the ability to auth both MAB and Dot1x together, eliminating delays while methods time out and fallback kicks in
- This gets the switch asking where should something go
- IBNS provides for auth failure, retries, even critical authentication

Automated ports – The magic

- ISE uses the internal endpoint database for MAB authentication
- Devices are imported via the ERS API
- Tagged with a “type” that maps back to our registration system

Automated ports – The magic

- ISE uses AD to learn Staff / Student groupings, as well as more specific groups for higher security rights
- This is applied to dot1x authentications in addition to machine type to provide 2 factors for network access
- Optional now, mandatory for staff access soon

Automated ports – The magic

- ISE Policy sets match these options against an authentication flow
- Once a match is found, the assigned profile and SGT is pushed

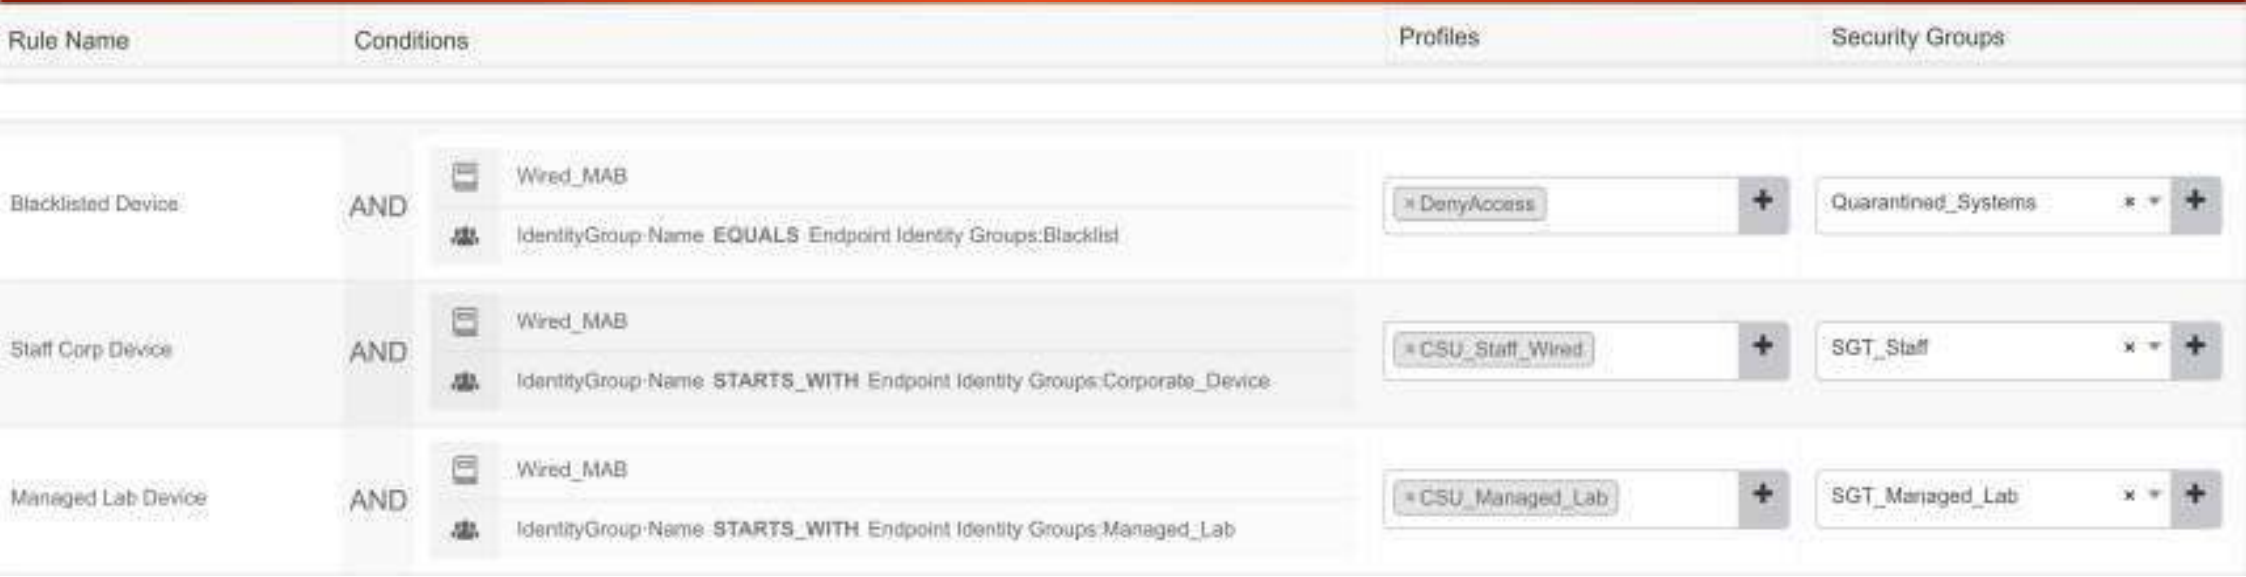

Putting it all together

- Switch provisioning setup handled via CSV into CSU integration
- Script takes switch member names, serial numbers, asset numbers and for 1st stack members IPv4 mgmt addresses

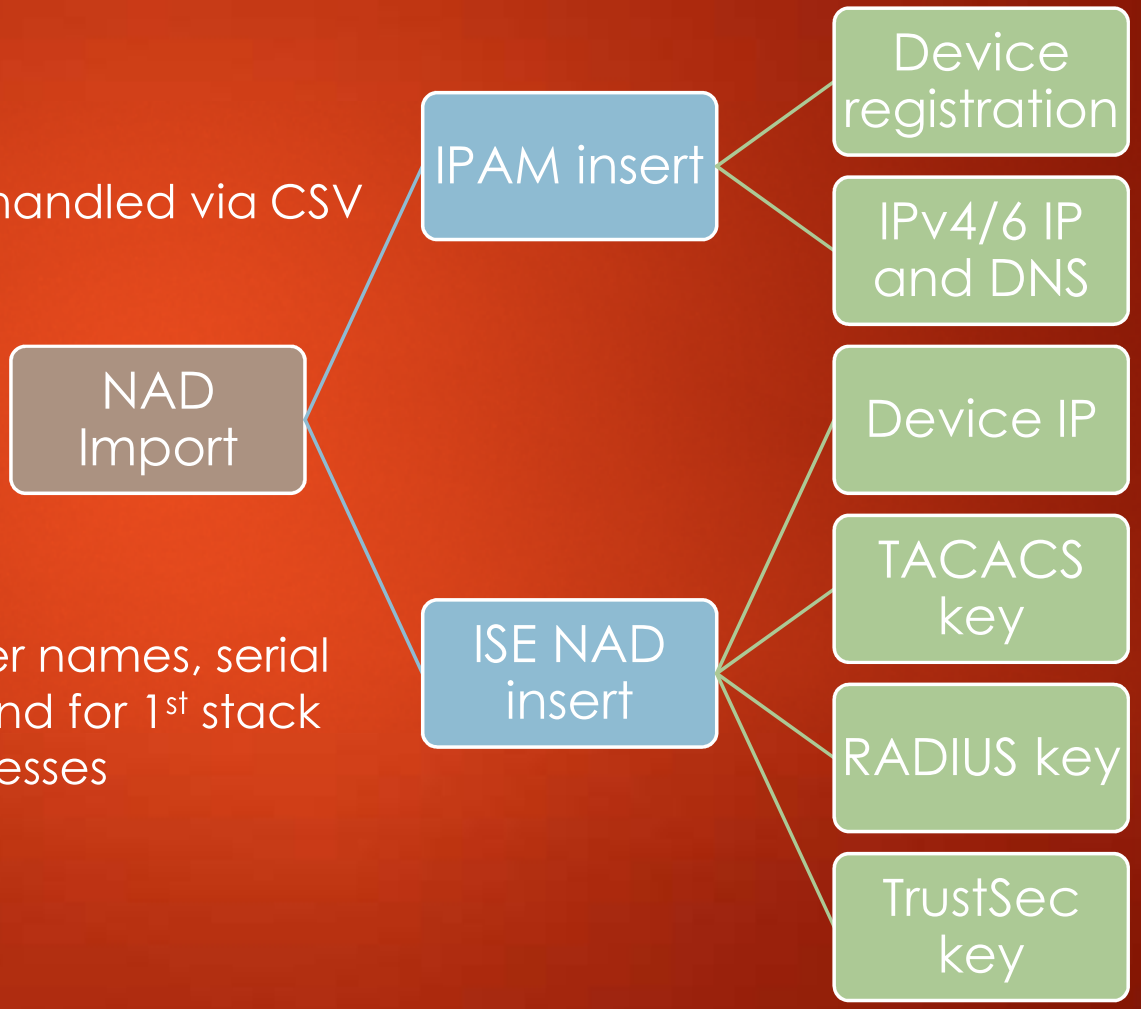

Putting it all together

- Host provisioning handled by our IPAM/Device management
- Entries are inserted / removed from ISE endpoint DB on demand
- CSU REST endpoint confirms if updates are needed, then hits ISE ERS
- A single "bulk load" was done to populate the DB initially

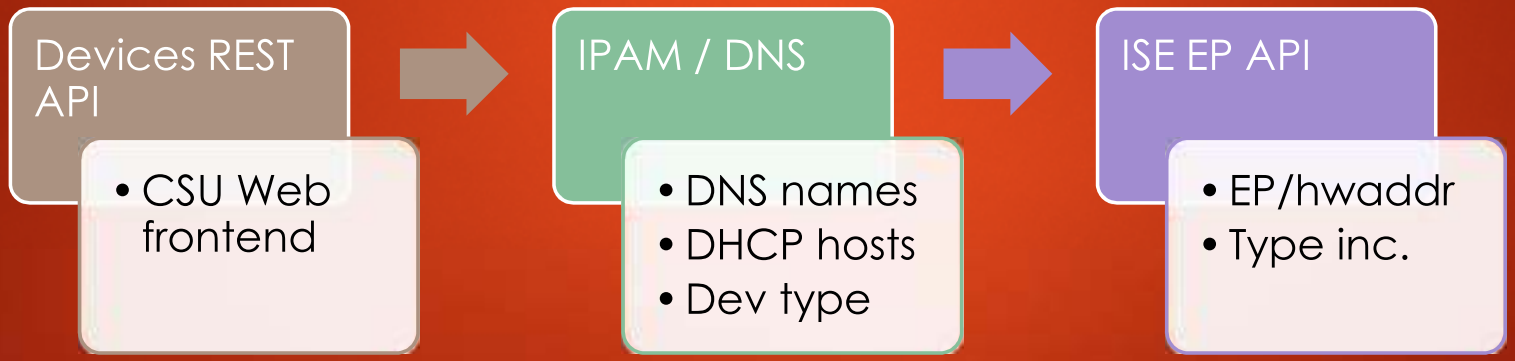

Putting it all together

- Switch replacements were started in October 2016 with Bathurst
- Each stack was taken offline during the day in fixed disruptions
- Swap-ins were done by an external contractor
- Patches were not required to be audited, devices were simply plugged back in wherever the cable could go
- WAPs were moved to Mgig ports where they weren't already

Upgrade project – Lessons learnt Models

- Mgig / Nbase-T has proven to be very reliable getting our APs running at 5gbps over some questionable UTP
- Mgig has a few caveats:
  - They don't go down to 10mbit, so that old BMS device may not work
  - You can't reliably port bundle over it – Speed adjustments will break it

Upgrade project – Lessons learnt Provisioning

- Set up your stacks with the same configuration on every member
- Pre-stage your new switches if possible for replacements

Upgrade project – Lessons learnt Systems

- Prime is useful, but generally a bit of a pain – We have HA problems every other week and a lot of what it does is not intuitive
- With that said, ISE is the exact opposite – It works flawlessly all the time, even during major upgrades and seems to be incredibly fault resistant

Upgrade project – Lessons learnt Authentications

- Load balancing switches auth by range is easy, but large single sources like a WLC is hard – **Use a proper HLB set up correctly**
- Devices that don't DHCP and don't speak unless spoken to (IoT sensors generally) will fail to be classified by the switch - DHCP everything!

Questions?