Experiences of mass CPE router management

Brian Candler

Inspired Gaming Group PLC
B.Candler@pobox.com

UKNOF10 – 21 May 2008

About INGG plc

- The leading provider of out-of-home networked entertainment systems, e.g.
  - “Itbox” in pubs
  - Digital jukeboxes
  - Fixed-odds betting terminals
  - Casino and bingo terminals
  - Large UK estate and rapid international growth
  - www.ingg.com

The UK ADSL network today

- Around 8,000 ADSL lines, L2TP delivery
- Assorted routers
  - Cisco 877W
  - Ericsson HN294
  - Ericsson ABS1000 (some)
  - Speedtouch ST5xx (few)
- Wholesale wi-fi service to The Cloud
- Assorted non-ADSL devices too
  - Compex WP54 bridges, Sarian HR4xxx GPRS

The problem: how do we manage all these?

- Generating unique (but similar) configs
- Remote config changes
  - e.g. wifi on/off, automatically triggered by business systems
  - e.g. change of SSID
- Remote firmware upgrades
  - in batches / en masse
- Tracking device presence

The original solution

- Originally we had only Ericsson HN294
- Using IP-IP tunnelling for The Cloud
- Basic requirements
  - “wifi on/off” triggered from business systems
  - remote firmware updates (lots of bugs!)
- Nothing suitable existed, so built our own

The original system

- Collection of Perl scripts, fork/exec
- Mostly CLI config, some web pages
- Event driven
  - RADIUS Accounting-Start packets were the triggers for pending actions
  - Reboot a router to start the process
  - “Vector” mapped state -> action, new state
- Static configs
  - all pre-generated and stored on disk

Architecture

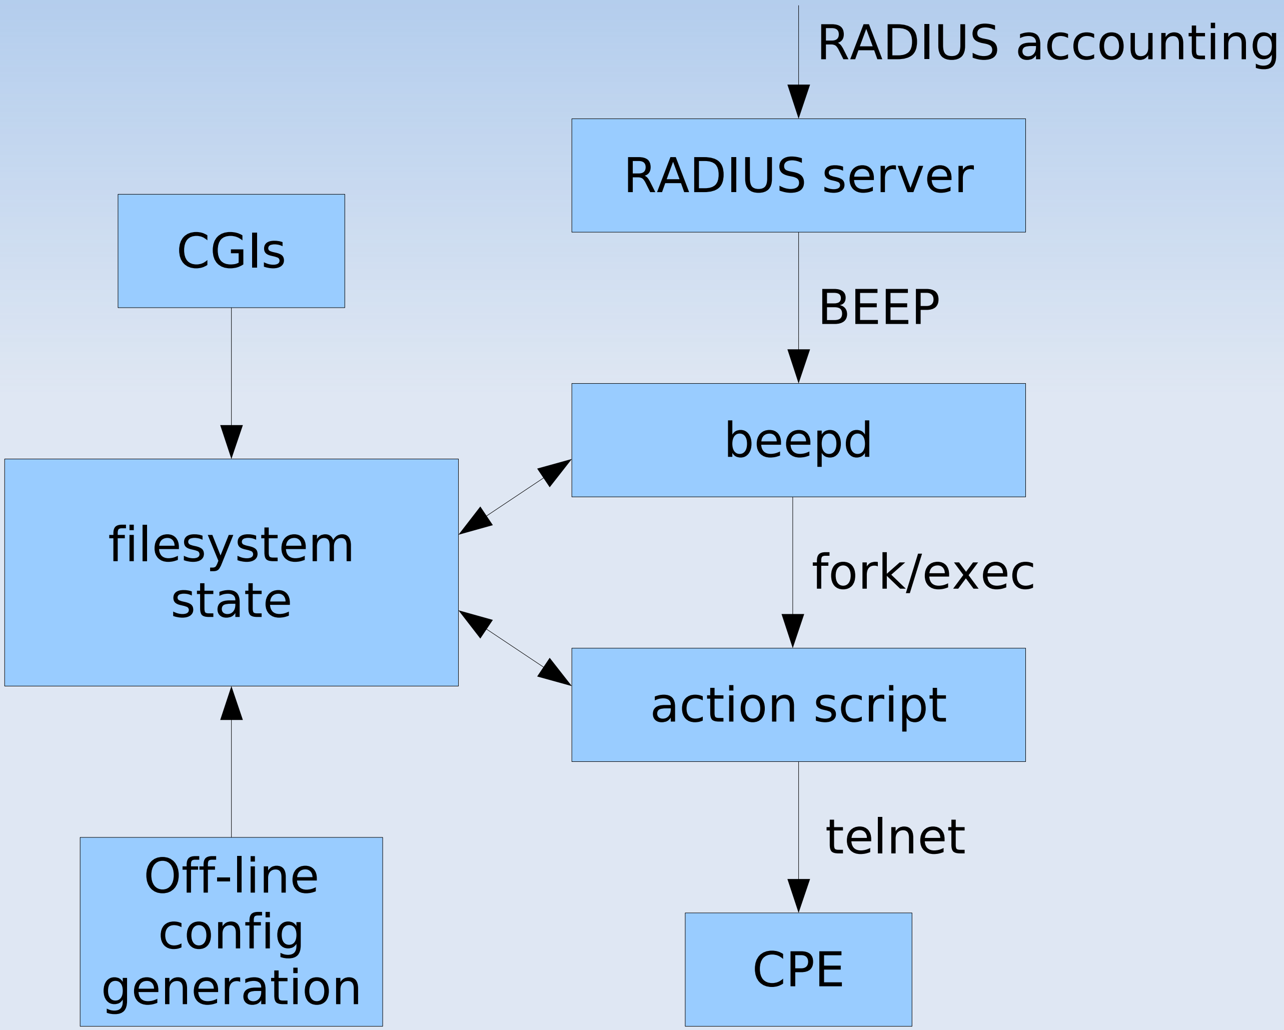

Looks simpler than it was :-)

- Did the job, but:
  - Hard to extend
  - Quite slow
  - RADIUS triggers only
- Relied on pre-generated configs on disk
  - one for wifi on, one for wifi off
  - doesn't scale, e.g. if you want to select 1 of 13 wifi channels
  - fixed subnet sizes

But lots of experience gained

- Reboot CPE before reconfiguring it
  - avoid problems with memory leaks
  - prove it comes up properly before changing it
  - and reboot to bring new config live
- Cheap home routers can be very fussy
  - sometimes had to downgrade from version Y to version X before upgrade to version Z
  - test, test, test!
  - ABS1000 (replacement for HN294) is a dud

The Next Generation

- Change was forced: HN294 went EOL
- Decided to move to Cisco 877W
  - “Enterprise” rather than “home” grade
  - IOS that we know and love (?)
  - Multi SSID; use a private SSID for wireless connectivity to our own devices; VRFs
  - TACACS
- Cisco's management platform (CNS)?
  - They were very coy about pricing
  - Complexity of integration

Meanwhile more devices...

- Compex WP54 wireless bridge
  - Itbox wireless client without the pain of Windows wireless networking
  - but we *need* a mass management platform
  - *dynamic* IP address via DHCP
- Patton Smartnode (VOIP), Sarian (3G)
- No mgmt platform which does all these
  - and none for the Compex at all
- So we ended up writing again

The new architecture

- Ruby on Rails for web interface
  - rapid, fluid development
- Ruby for “back end” processes
  - reasonable libraries for Telnet, SSH
  - but had to write own TFTP
- Oracle for state storage
  - we already had an Oracle cluster
- Subversion for config templates
  - check out on multiple servers

Architecture

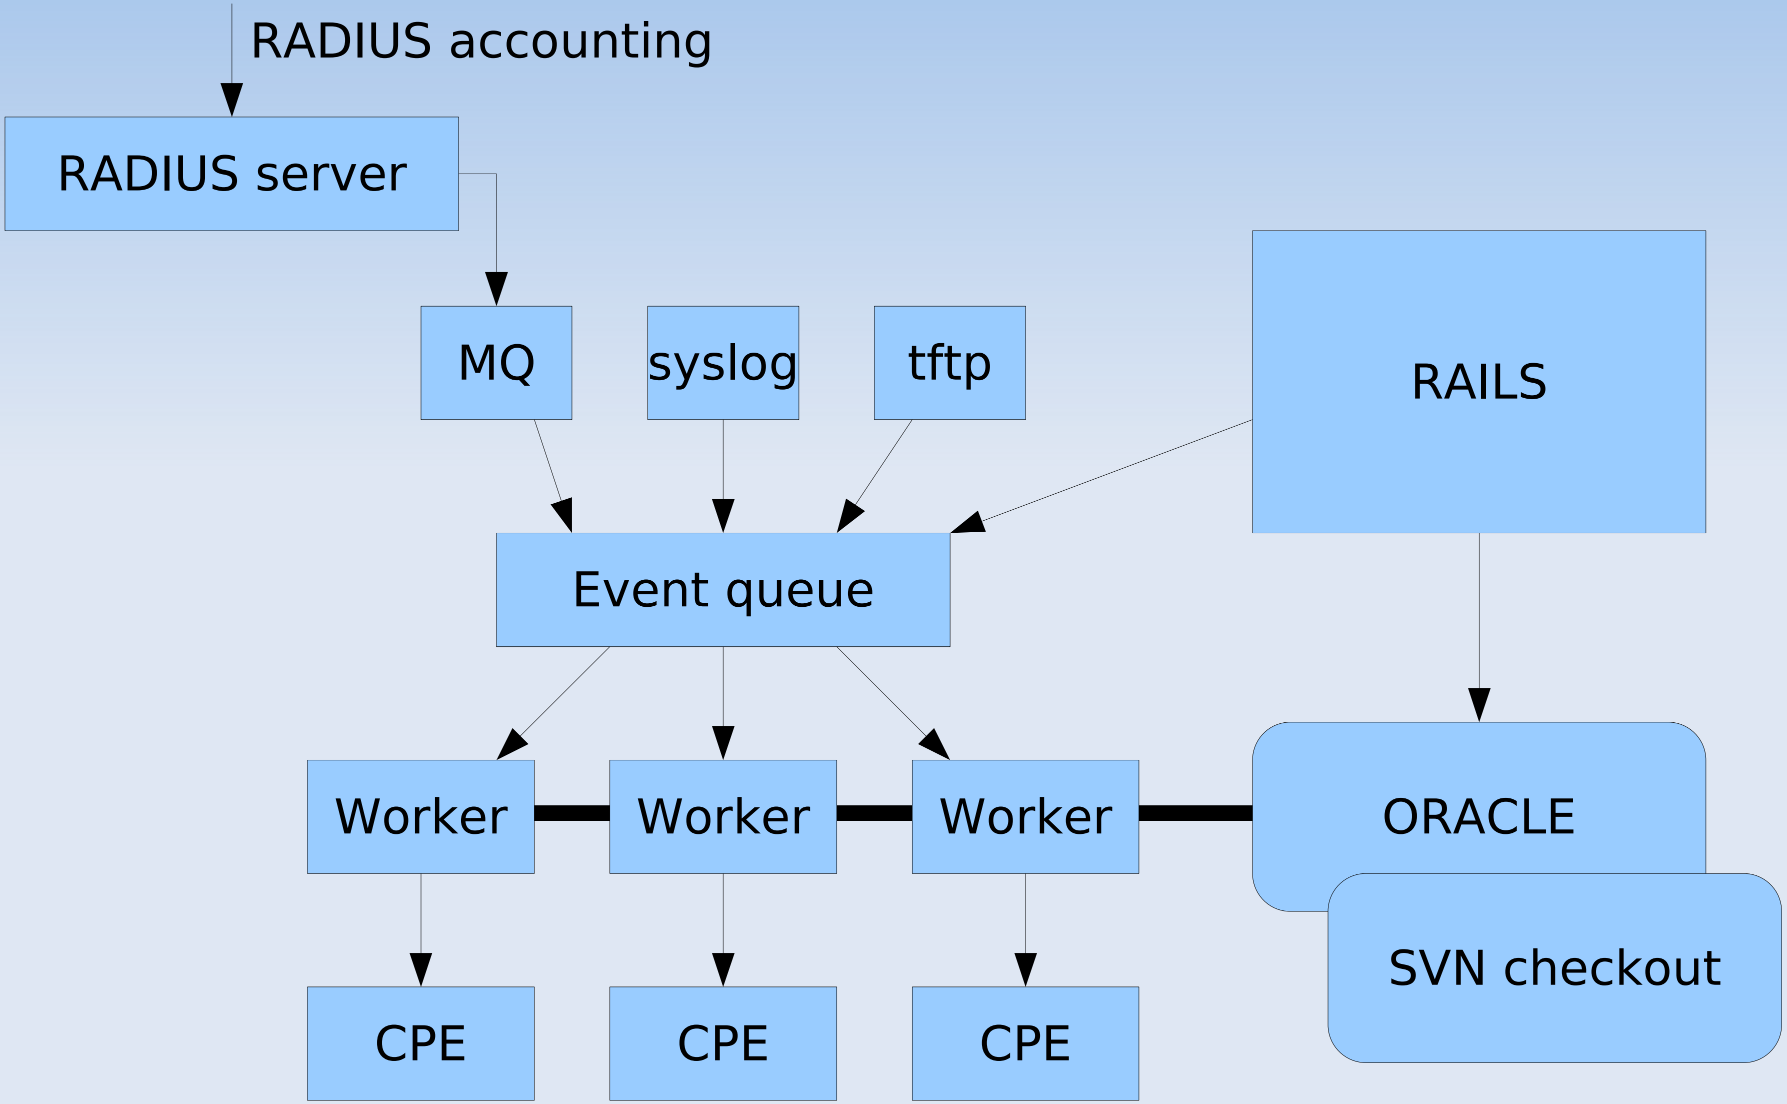

Persistent worker processes

- Take jobs from a (memory) queue
- Single thread: handle one event/one device connection to completion
  - No concurrency issues
- Multiple processes
- Decrypts master password file at startup
  - Oracle doesn't store actual passwords; only labels like “NETENTLOGIN-1”
- Integrated Ruby TFTP+HTTP servers

New features

- Generate device configs on the fly
  - erb template plus db “settings” (attr=val)
  - Bundle of default settings: “target config”
  - Settings can be overridden for individual devices (e.g. wifi_enable, wifi_chan)
- Allocations assigned on demand
- Can also perform bulk assignments; download batch of configs as a zip file
- Generates RADIUS and IP-IP config files

Example settings and template

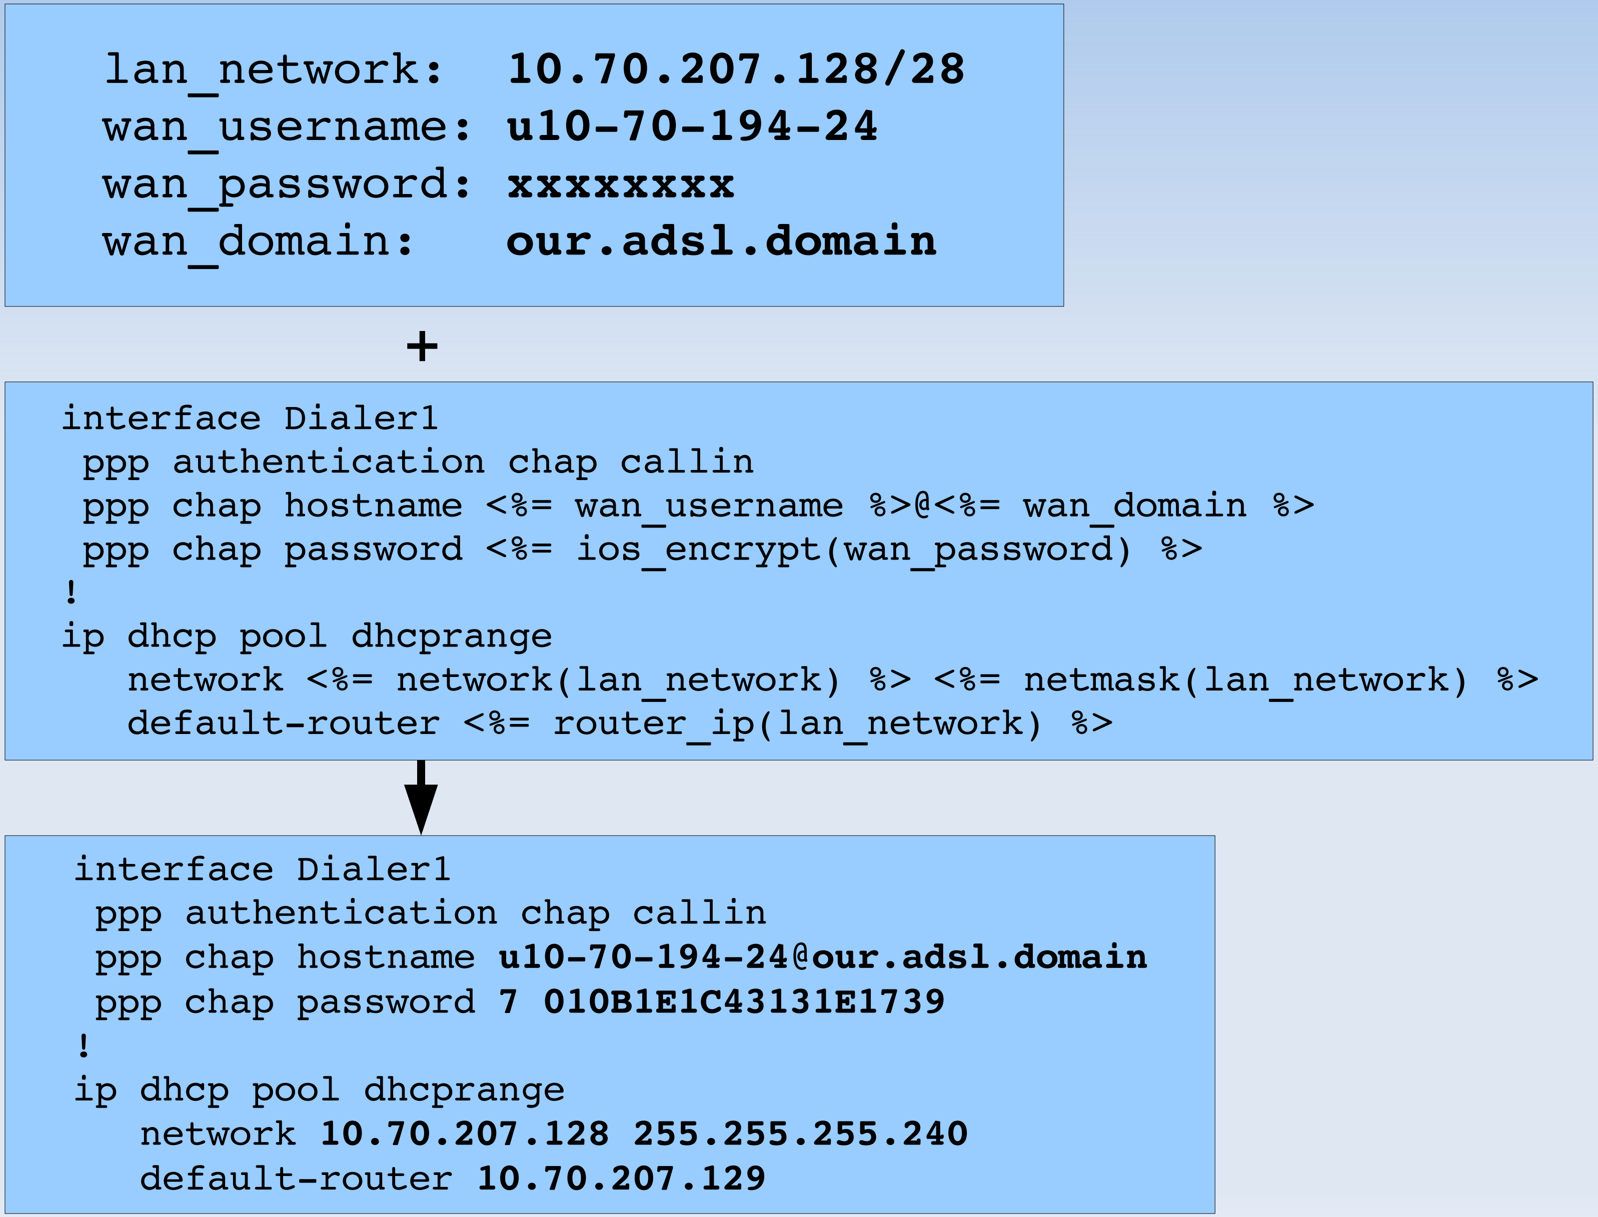

```
lan_network:  10.70.207.128/28
wan_username: u10-70-194-24
wan_password: xxxxxxxx
wan_domain:   our.adsl.domain
```

+

```
interface Dialer1
 ppp authentication chap callin
 ppp chap hostname <%= wan_username %>@<%= wan_domain %>
 ppp chap password <%= ios_encrypt(wan_password) %>
!
ip dhcp pool dhcprange
   network <%= network(lan_network) %> <%= netmask(lan_network) %>
   default-router <%= router_ip(lan_network) %>
```

```
interface Dialer1
 ppp authentication chap callin
 ppp chap hostname u10-70-194-24@our.adsl.domain
 ppp chap password 7 010B1E1C43131E1739
!
ip dhcp pool dhcprange
   network 10.70.207.128 255.255.255.240
   default-router 10.70.207.129
```

Devices have “functions”

- We ended up using the 877W for lots of different purposes
  - standard ADSL router
  - Internet VPN router
  - VPN+wireless router... etc
- Setting the “function” limits what configs can be applied
  - avoid accidentally applying the wrong type of config to a device!

Device drivers

- Drivers provide a simple device abstraction, e.g.
  - “read your MAC address”
  - “upload this firmware”
  - “upload this config”
  - “download this config”
  - “reboot”
- Integrate a new device in 1-2 days
  - some Ruby magic helps (e.g. mount a ZIP file and serve it via TFTP, for Patton)

Vectors (action scripts)

- What to do when a given event type arrives, when we are in a given state
- Invoke device driver, perform some work, change state (i.e. wait for next event)
- “Generic” vector works for most things
  - download config and compare to target
  - reboot if uptime too large
  - upload firmware if required
  - upload config, reboot
  - download config and compare again

Dynamic IP management

- Incoming events have associated IP address, e.g.
  - Cisco: RADIUS Framed-IP-Address
  - Compex: syslog source IP address
- Event data locates the Oracle record
  - RADIUS User-Name
  - syslog message contains MAC address
- Store updated IP address in DB
- Then trigger pending action, if any

Probing

- What happens when we see a new device (unknown RADIUS username or MAC?)
  - Check its telnet and ssh banners against known patterns
  - Try logging in with known passwords
  - Register the device in Oracle
  - Assumes only “trusted” devices are connected!

Desktop proxy

- Small Openwrt (Linux) router
  - e.g. Buffalo WHR54GS (£25)
- Scripts configure it as a NAT relay between LAN and a local port
  - acts as both DHCP server and client, to auto-detect the other side
- Connect a device on its factory-default IP (e.g. 10.10.10.1), makes it reachable from the central system
- Cool :-)

But the system is still...

- Event driven
  - RADIUS accounting (via MQ not BEEP)
  - syslog (when Compex obtains DHCP lease)
  - tftp (periodic TFTP request from Patton)
  - operator triggers via web interface
- “Push and reboot” operation
  - copy tftp startup-config; reload
  - however this logic is easily changed with a new vector

Limitations

- Files in Subversion
  - adding a template means check out, modify, check in, and check out on all servers
- IP management
  - IP allocations still prepared off-line using the Perl code, then imported into Oracle
  - importing, say, a new /16 as a bunch of /27's involves a bit of specialist work
- Probes Ericssons and Speedtouch, but doesn't manage them

Experience gained: Rails

- Rails production environment takes a bit more effort than expected
  - New book “Deploying Rails Applications” helps (pragprog.com)
  - Need a proxy which sends only 1 HTTP request at a time to each mongrel process
    - We use Apache 2.0 in front of pen
    - Nginx plus fair proxy module is a new alternative
  - Set up monit to start/stop all processes
  - Use capistrano intelligently
    - Separate deployment from development

Experience gained: Cisco CPE

- Some CLI differences between versions
  - Should we use SNMP set instead??
- 877W has only 24MB flash
  - IOS image is 17MB
  - Must erase old image before installing new
  - Remote reflashing is very dangerous
- Newly-announced 88X range has 128MB
  - however there is SDSL model but no ADSL !!
- Problems managing vlan.dat

Experience gained: Misc

- TFTP is very slow over ADSL links
  - No “window” - req1 ack1 req2 ack2...
  - 512 byte packets (unless negotiated up)
  - 50ms RTT => 20 pps => 10KB/s
- Saw problems with Cisco using HTTP to transfer firmware images
  - Transfer aborted but CLI said success
  - Could use 'verify /md5' to check upload

Vendors and management

- Inconsistent approaches
  - some only have web GUI
    - Compex added TFTP via CLI for us, thanks!
  - some “pull” configs from server
  - some “push” configs to device
  - mix of tftp/ftp/http, very little sftp/scp
  - Cisco CNS (closed and proprietary)
  - Anyone for TR-069?? Very SOAPy

Future developments? Questions?