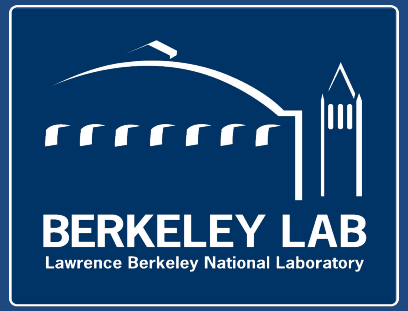

U.S. DEPARTMENT OF
ENERGY

UNIVERSITY OF
CALIFORNIA

CI Engineering Lunch & Learn Series

LBNL's Cloudflare Lessons Learned

Michael Smitasin, Cyber Security Engineer

cloudflare@lbl.gov

2020-10-23

Quick intro (1)

- Me
  - 7 years at LBNL
    - 2.5 years in Cyber Security
    - 4.5 years in Network Engineering
  - SCinet 2015-2017
  - IT of some form since ~2005

- This talk
  - Not a sales pitch, I don't work for Cloudflare (or own CF stock)
  - Doesn't cover all their services, just what we're using

Quick intro (2)

- Department of Energy national lab, but "like a research university"

- LBNL has ~545 websites using Cloudflare as of Sep 2020
  - Including ~125 Business Systems / Enterprise sites
- Remaining:
  - ~900 on-prem
  - ~200 cloud or off-prem

- First "production" site moved **2018-04-27**

- www.lbl.gov moved **2018-10-17**

Outline

- Laying the Groundwork
- Our specific use
- Day-to-Day operations
- Architecture
- Considerations before implementing
- Known Limitations
- Lessons Learned
- Overall experience

Laying the Groundwork (1)

- Motivations
  - Actual damage
    - past web attacks, alarming defacements
  - Reputational damage
    - misinterpretation of low value web attacks ("cyclotron compromised")

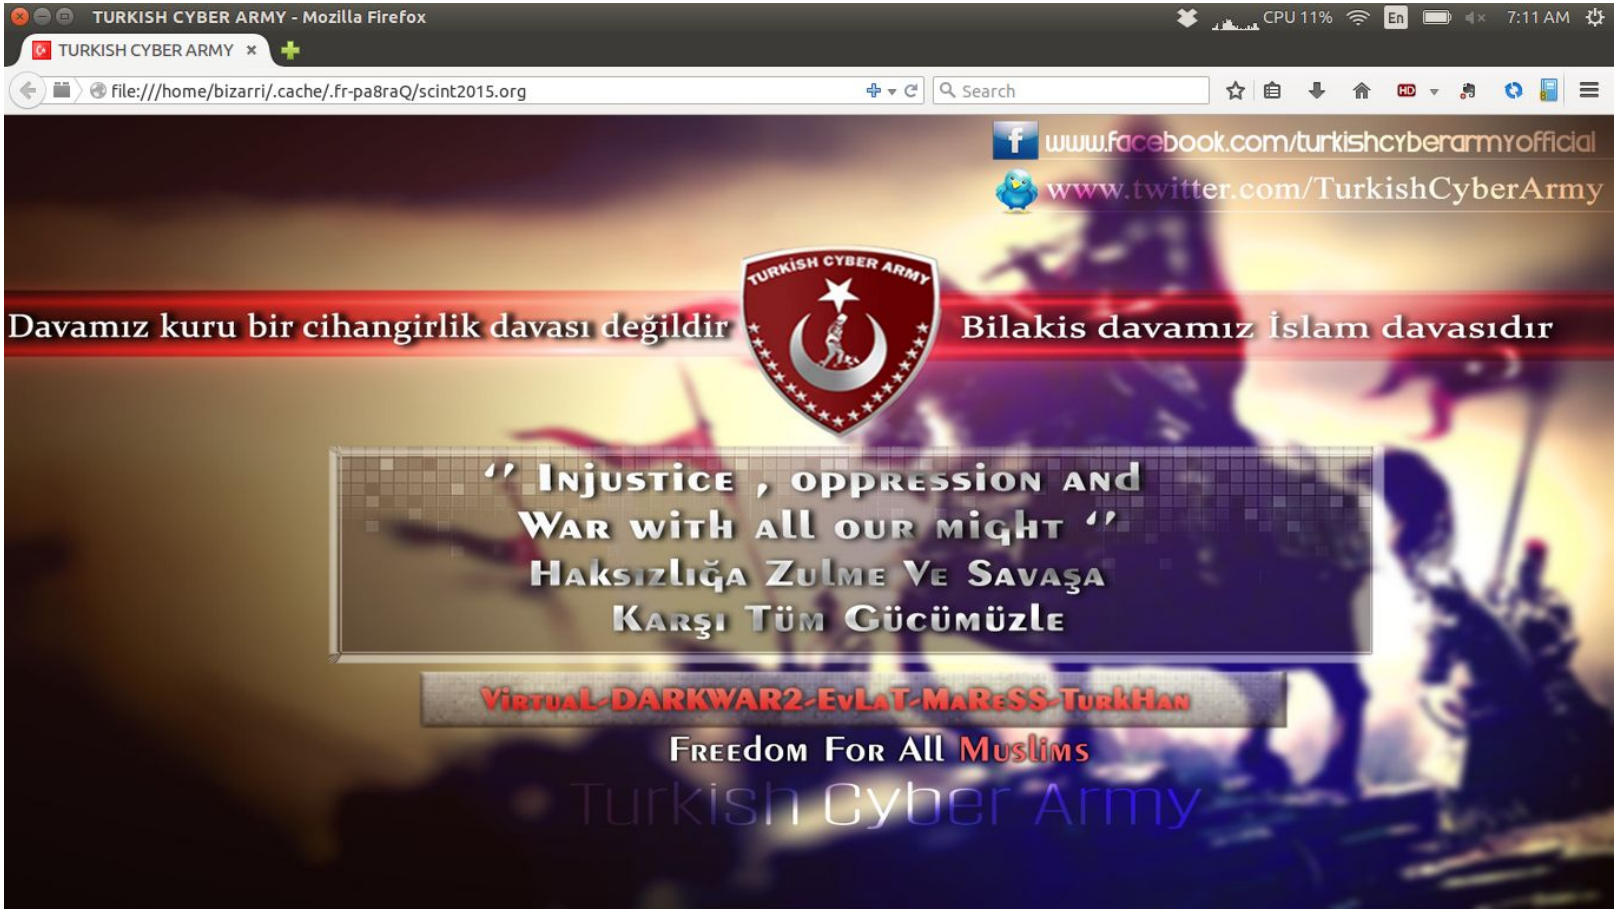
TURKISH CYBER ARMY - Mozilla Firefox
TURKISH CYBER ARMY
file:///home/bizarri/.cache/.fr-pa8raQ/scint2015.org
Search
www.facebook.com/turkishcyberarmyofficial
www.twitter.com/TurkishCyberArmy
Davamız kuru bir cihangirlik davası değildir
Bilakis davamız İslam davasıdır
TURKISH CYBER ARMY
'' Injustice , oppression and
War with all our might ''
Haksızlığa Zulme Ve Savaşa
Karşı Tüm Gücümüzle
VirtuaL-DARKWAR2-EvLaT-MaReSS-TurkHan
Freedom For All Muslims
Turkish Cyber Army

U.S. NEWS

U.S. Infrastructure Can Be Hacked With Google, Simple Passwords

April 3, 2016, 7:07 AM PDT / Updated April 3, 2016, 7:07 AM PDT

By Chris Francescani

Authorities say the Iranian computer hack of a New York dam is the symptom of a huge weakness in the U.S. infrastructure -- dams, stadiums, traffic controls and power grids that can be accessed by anyone, including hostile nations or terrorists -- with simple passwords or no passwords at all.

Who Is Vulnerable?

The threat of cyber-attacks spans every sector of the U.S. economy, experts said.

In recent years, independent "white-hat" security researchers have shown they can access cities' traffic control systems and license plate reader networks, sports stadiums, car washes, a hockey rink in Denmark, a Texas water plant, the particle-accelerating cyclotron at the Lawrence Berkeley National Laboratory, even an Olympic arena.

Laying the Groundwork (2)

- Lack of visibility with off-prem sites
  - traditional VMs... but also SaaS, serverless apps, etc
  - CF = log visibility into off-prem sites (like Apache ssl.log)
  - Not quite SSL-decryption (no full pcaps)

- 'cause DOE says so: Binding Operational Directive 18-01

- New tech - Web Application Firewall (WAF) protection
  - Alternatives we looked at
    - F5, Imperva, VPN tunnels, web server log forwarding to syslog

Architecture

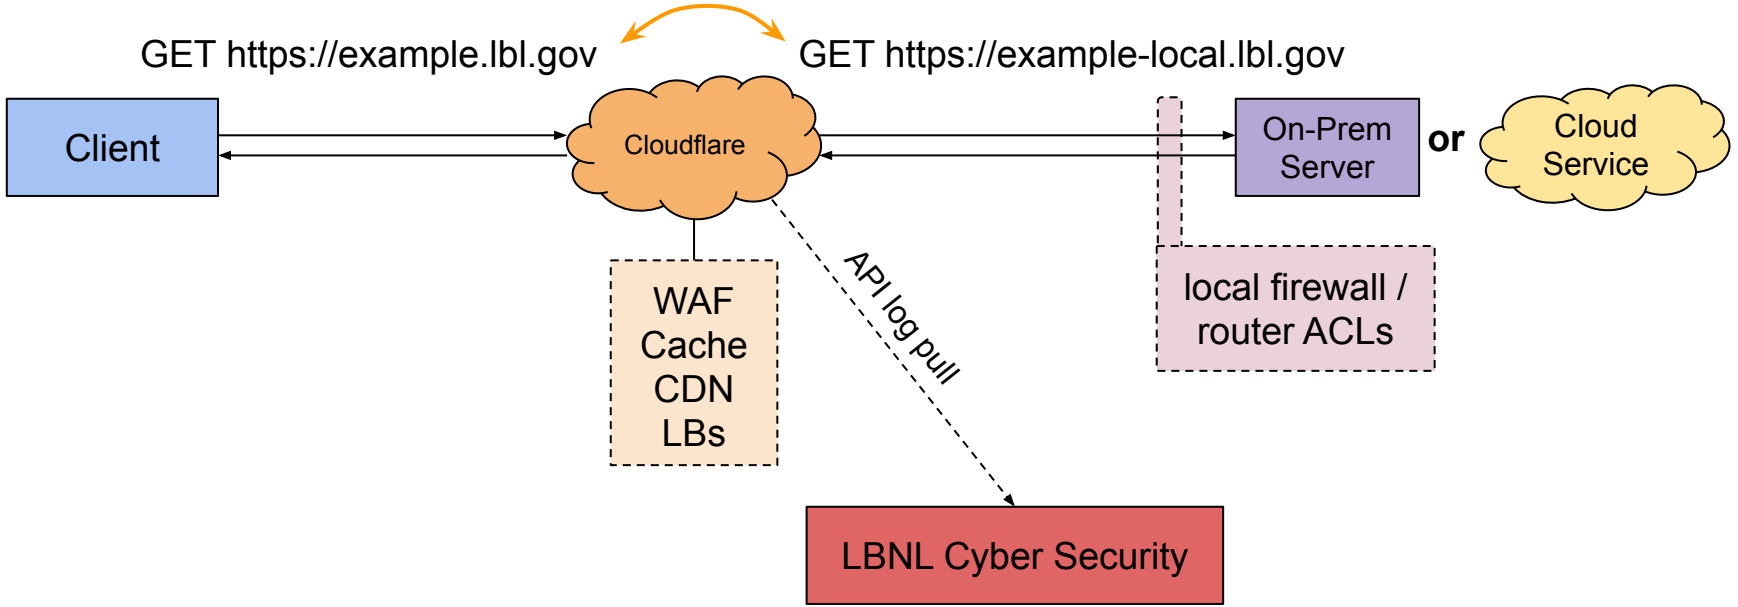

Our Specific Use (1)

- Reverse Proxy for HTTP/HTTPS only
- Still run authoritative DNS in-house (BIND), use our own registrar(s)

- Web Application Firewall (+ IP-based restrictions)
- SSL/TLS certs and other BOD 18-01 stuff for front-end
- Cache/CDN/Load-Balancer
- Logpull of HTTP/S logs + Cloudflare audit logs via API queries

- Future:
  - Cloudflare "Access" as ~pseudo-web-VPN
  - Spectrum (non-HTTP proxy), but concerns about vendor-lock

30-day traffic stats

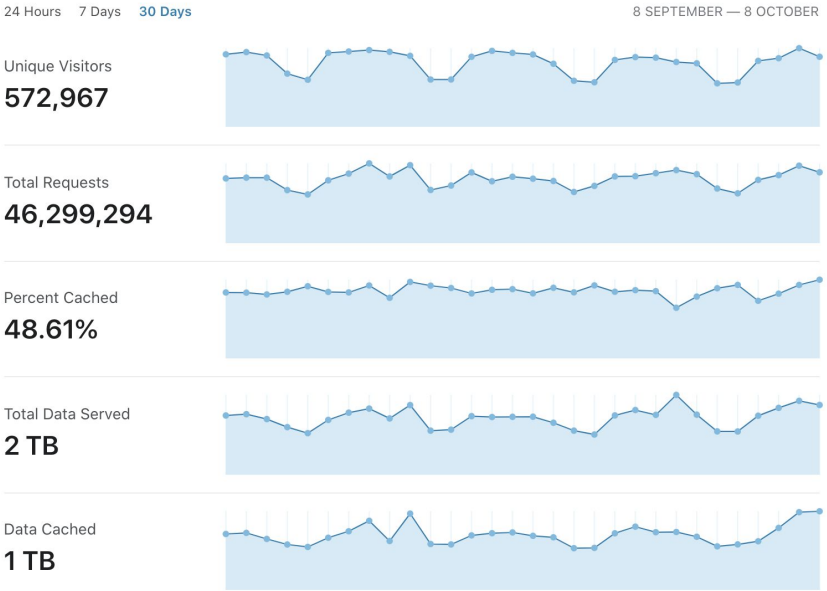

72 hr WAF stats

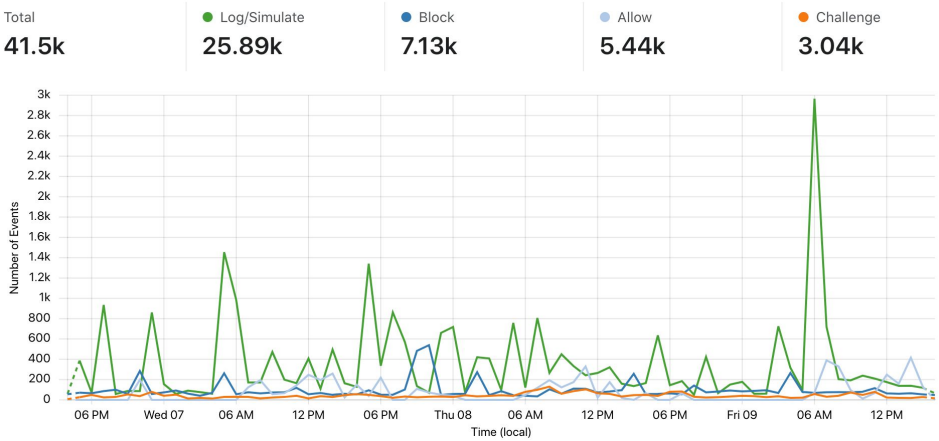

Our Specific Use (2)

- Example WAF drop event from API logpull:


```
{
  "CacheCacheStatus": "unknown",
  "ClientIP": "192.0.2.66",
  "ClientRequestBytes": 2119,
  "ClientRequestHost": "www.lbl.gov",
  "ClientRequestMethod": "GET",
  "ClientRequestProtocol": "HTTP/1.1",
  "ClientRequestReferer": "",
  "ClientRequestURI":
"//laravel/vendor/phpunit/phpunit/src/Util/PHP/eval-stdin.php",
  "ClientRequestUserAgent": "python-requests/2.24.0",
  "ClientSSLCipher": "ECDHE-ECDSA-AES128-GCM-SHA256",
  "ClientSSLProtocol": "TLSv1.2",
  "ClientSrcPort": 50663,
  "EdgeServerIP": "",
  "EdgeStartTimestamp": 1601016787862000000,
  "OriginIP": "",
  "OriginResponseBytes": 0,
  "OriginResponseStatus": 0,
  "WAFAction": "drop",
  "WAFFlags": "0",
  "WAFMatchedVar": "",
  "WAFProfile": "med",
  "WAFRuleID": "PHP100012",
  "WAFRuleMessage": "PHP, PHPUnit - Code Injection -
CVE:CVE-2017-9841"
}
```

Our Specific Use (3)

Default Error Page

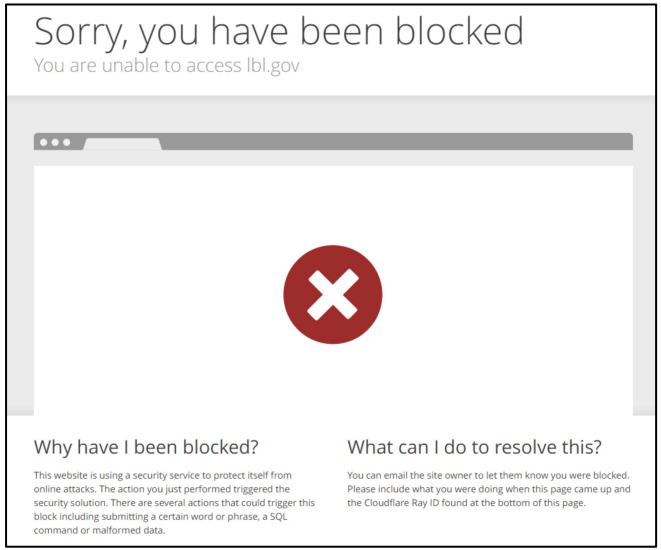

Custom Error Page

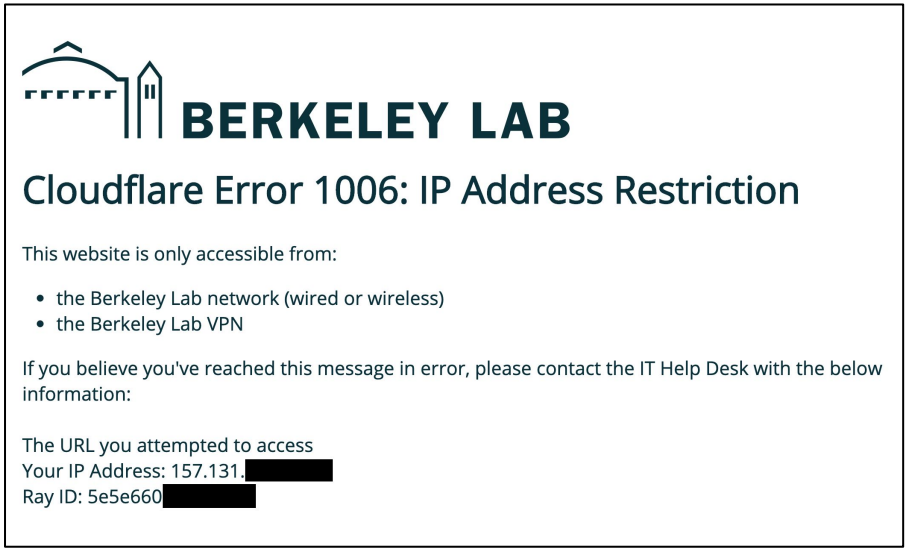

Day-to-Day operations (1)

- Add something to use Cloudflare
    - **Cloudflare:** example1.lbl.gov -> example1-local.lbl.gov
      (generates example1.lbl.gov.cdn.cloudflare.net)
    - **DNS:** example1.lbl.gov CNAME example1.lbl.gov.cdn.cloudflare.net
    - **Border ACLs:** add origin server IP address to prefix-list

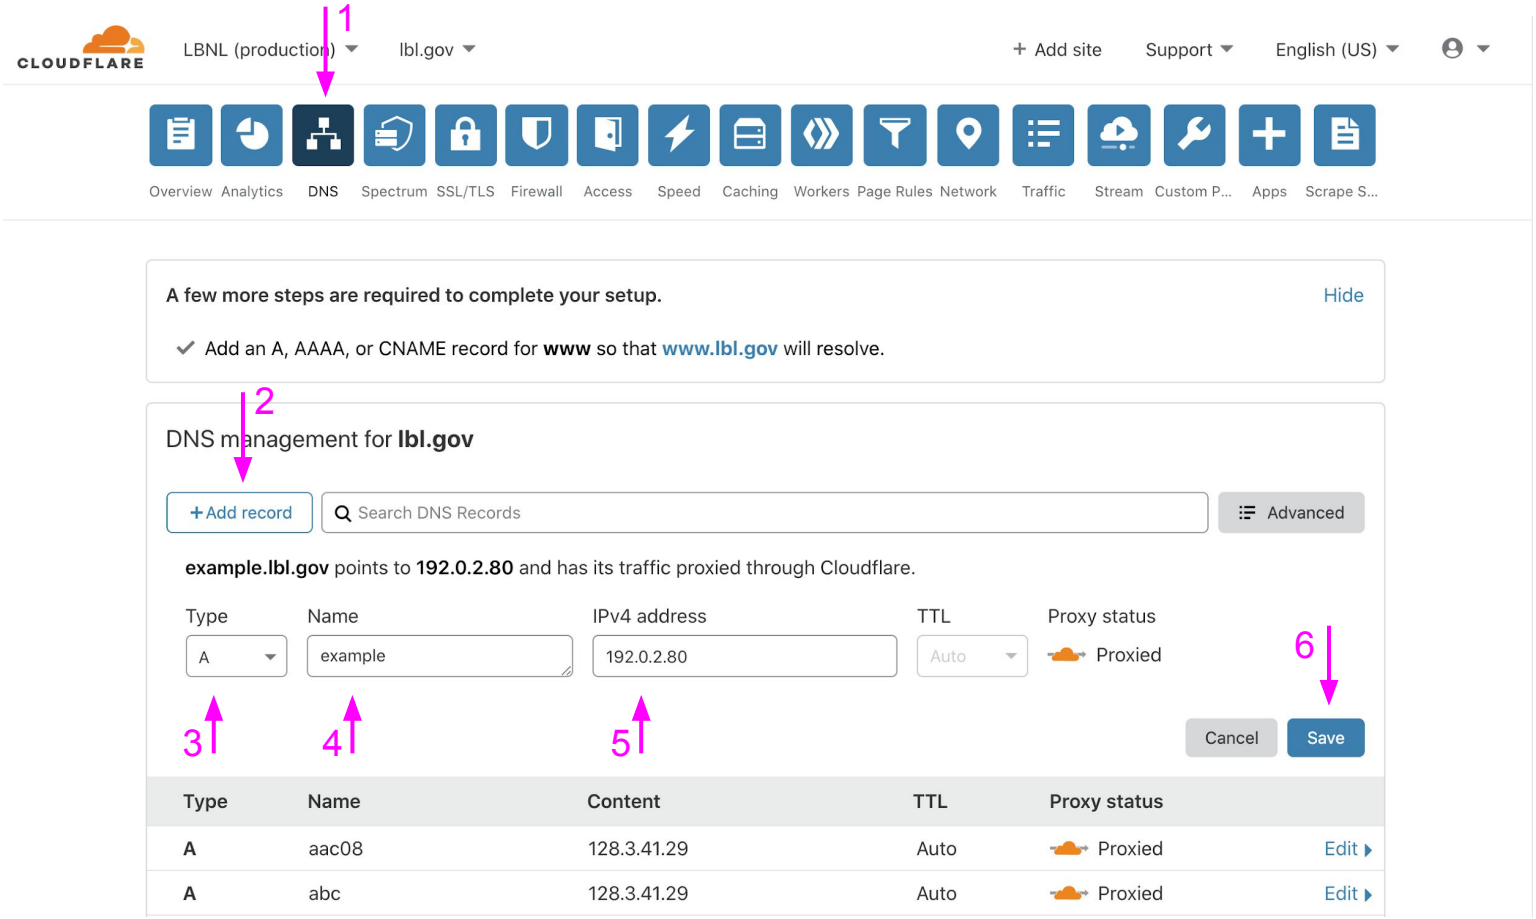
CLOUDFLARE
LBNL (production)
lbl.gov
+ Add site
Support
English (US)
1
Overview
Analytics
DNS
Spectrum
SSL/TLS
Firewall
Access
Speed
Caching
Workers
Page Rules
Network
Traffic
Stream
Custom P...
Apps
Scrape S...
A few more steps are required to complete your setup.
Hide
Add an A, AAAA, or CNAME record for www so that www.lbl.gov will resolve.
2
DNS management for lbl.gov
+ Add record
Search DNS Records
Advanced
example.lbl.gov points to 192.0.2.80 and has its traffic proxied through Cloudflare.
Type
Name
IPv4 address
TTL
Proxy status
A
example
192.0.2.80
Auto
Proxied
6
3
4
5
Cancel
Save
Type
Name
Content
TTL
Proxy status
A
aac08
128.3.41.29
Auto
Proxied
Edit
A
abc
128.3.41.29
Auto
Proxied
Edit

Day-to-Day operations (2)

- Tune some WAF setting due to False Positives
  - Check the logs / WAF events
  - Identify responsible rule
  - Change action on rule or disable WAF for URI ("Page Rule")

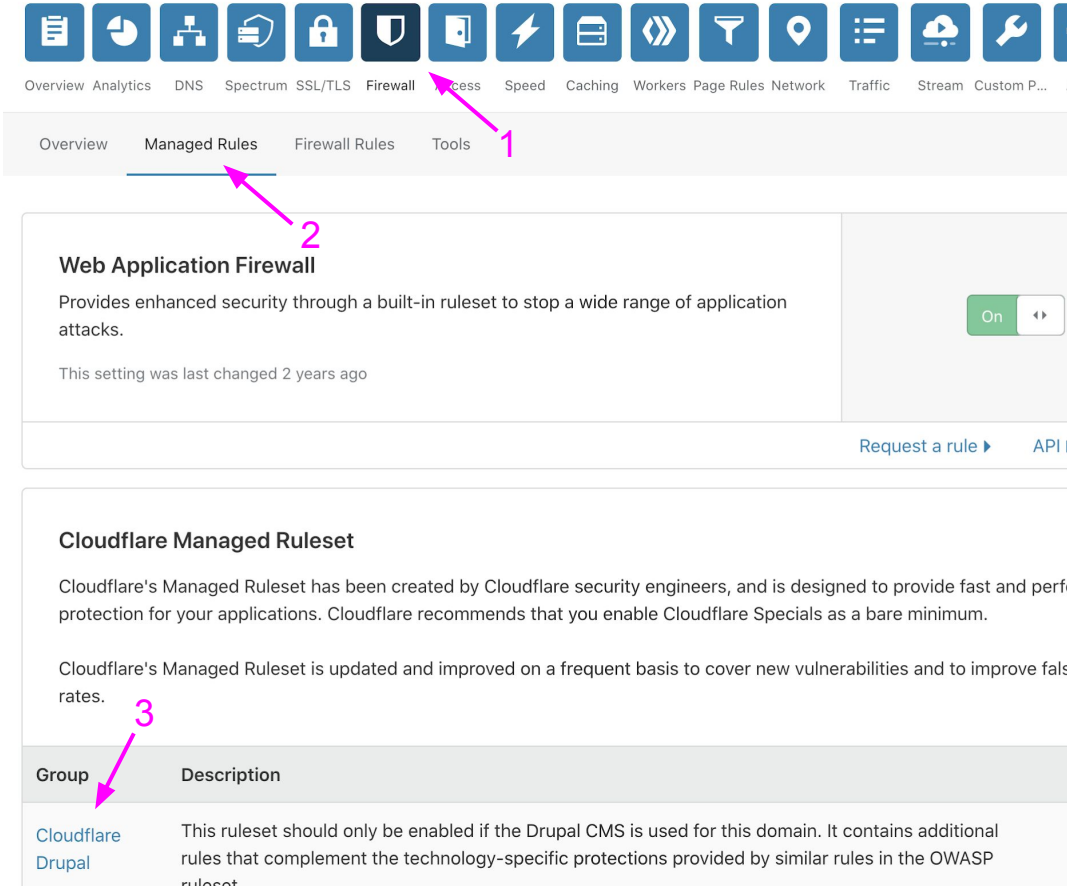
Overview
Analytics
DNS
Spectrum
SSL/TLS
Firewall
Access
Speed
Caching
Workers
Page Rules
Network
Traffic
Stream
Custom P...
1
Overview
Managed Rules
Firewall Rules
Tools
2
Web Application Firewall
Provides enhanced security through a built-in ruleset to stop a wide range of application attacks.
This setting was last changed 2 years ago
On
Request a rule
API
Cloudflare Managed Ruleset
Cloudflare's Managed Ruleset has been created by Cloudflare security engineers, and is designed to provide fast and perf
protection for your applications. Cloudflare recommends that you enable Cloudflare Specials as a bare minimum.
Cloudflare's Managed Ruleset is updated and improved on a frequent basis to cover new vulnerabilities and to improve fals
rates.
3
Group
Description
Cloudflare Drupal
This ruleset should only be enabled if the Drupal CMS is used for this domain. It contains additional rules that complement the technology-specific protections provided by similar rules in the OWASP

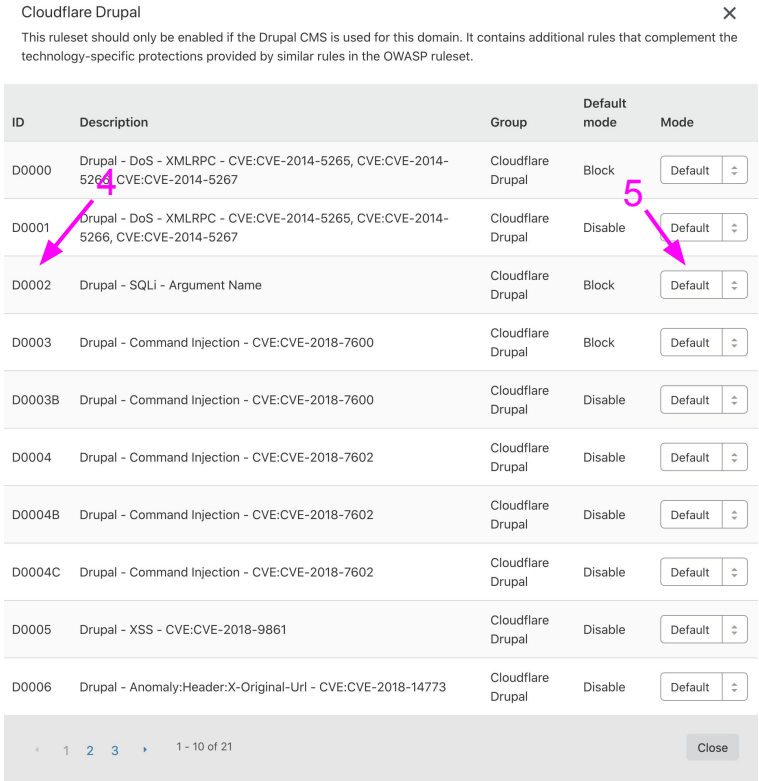
Cloudflare Drupal
This ruleset should only be enabled if the Drupal CMS is used for this domain. It contains additional rules that complement the technology-specific protections provided by similar rules in the OWASP ruleset.
4
5
ID
Description
Group
Default mode
Mode
D0000
Drupal - DoS - XMLRPC - CVE:CVE-2014-5265, CVE:CVE-2014-5266, CVE:CVE-2014-5267
Cloudflare Drupal
Block
Default
D0001
Drupal - DoS - XMLRPC - CVE:CVE-2014-5265, CVE:CVE-2014-5266, CVE:CVE-2014-5267
Cloudflare Drupal
Disable
Default
D0002
Drupal - SQLi - Argument Name
Cloudflare Drupal
Block
Default
D0003
Drupal - Command Injection - CVE:CVE-2018-7600
Cloudflare Drupal
Block
Default
D0003B
Drupal - Command Injection - CVE:CVE-2018-7600
Cloudflare Drupal
Disable
Default
D0004
Drupal - Command Injection - CVE:CVE-2018-7602
Cloudflare Drupal
Disable
Default
D0004B
Drupal - Command Injection - CVE:CVE-2018-7602
Cloudflare Drupal
Disable
Default
D0004C
Drupal - Command Injection - CVE:CVE-2018-7602
Cloudflare Drupal
Disable
Default
D0005
Drupal - XSS - CVE:CVE-2018-9861
Cloudflare Drupal
Disable
Default
D0006
Drupal - Anomaly:Header:X-Original-Url - CVE:CVE-2018-14773
Cloudflare Drupal
Disable
Default
1 2 3
1 - 10 of 21
Close

Considerations before implementing

- Authoritative DNS moves to Cloudflare?
- Preserving local or management DNS records
  - We require DNS records for static IPs
- Handling False Positives / "Unintended Blocks"
- Snowflakes vs. Monolithic Configuration
  - Config management? Backups? Change control?
- Centralized vs. Distributed Control
- Visibility
  - (how does anyone but CF admins know what's going on?)
- What gets put in Cloudflare? What doesn't?
  - Probably not: PerfSonar, Globus Endpoints, etc

Known Limitations (1)

- Possible to bypass Cloudflare by going directly to the Origin Server IP
  - Lock down Origin Server to CF IPs - https://www.cloudflare.com/ips/

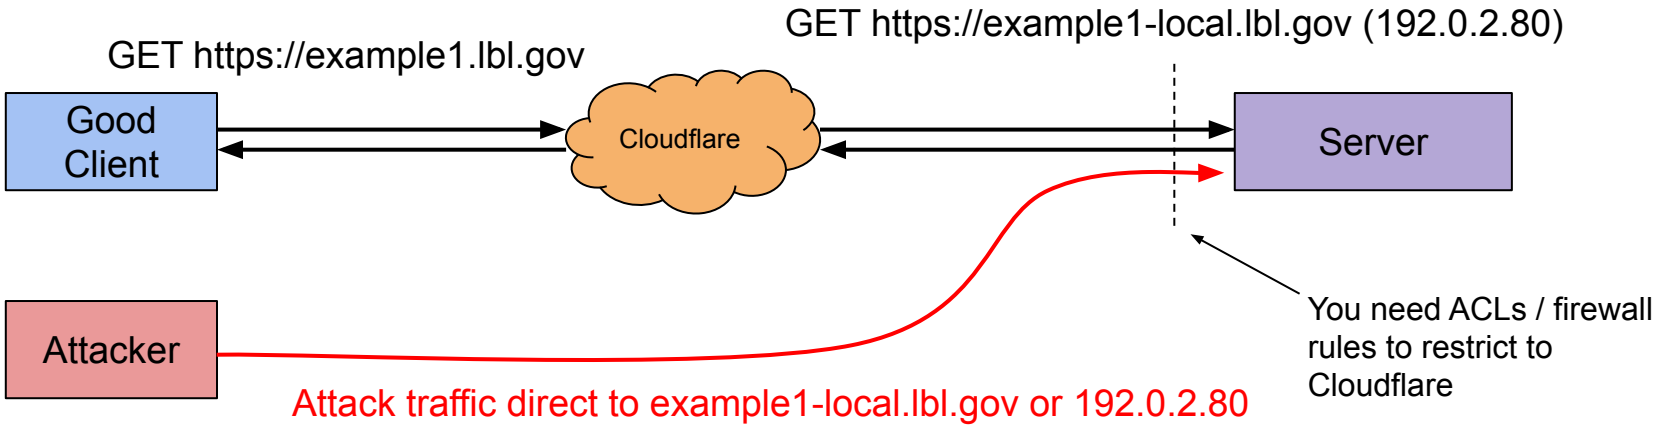

Known Limitations (2)

- Most rich functionality is HTTP/HTTPS
  - Spectrum = non-HTTP proxy, but lacking parity

- You may need to whitelist Cloudflare IPs in your IDS (could interpret a Cloudflare IP making lots of connections to multiple Origin Servers as attacks)
  - https://www.cloudflare.com/ips/

- No rate limiting by default (paid feature?), the Cache/CDN will only do so much, Origin Server can still get overwhelmed by "legitimate" traffic

Known Limitations (3)

- Structure - permissions/role granularity, "Zone Wide" settings, splitting out sites

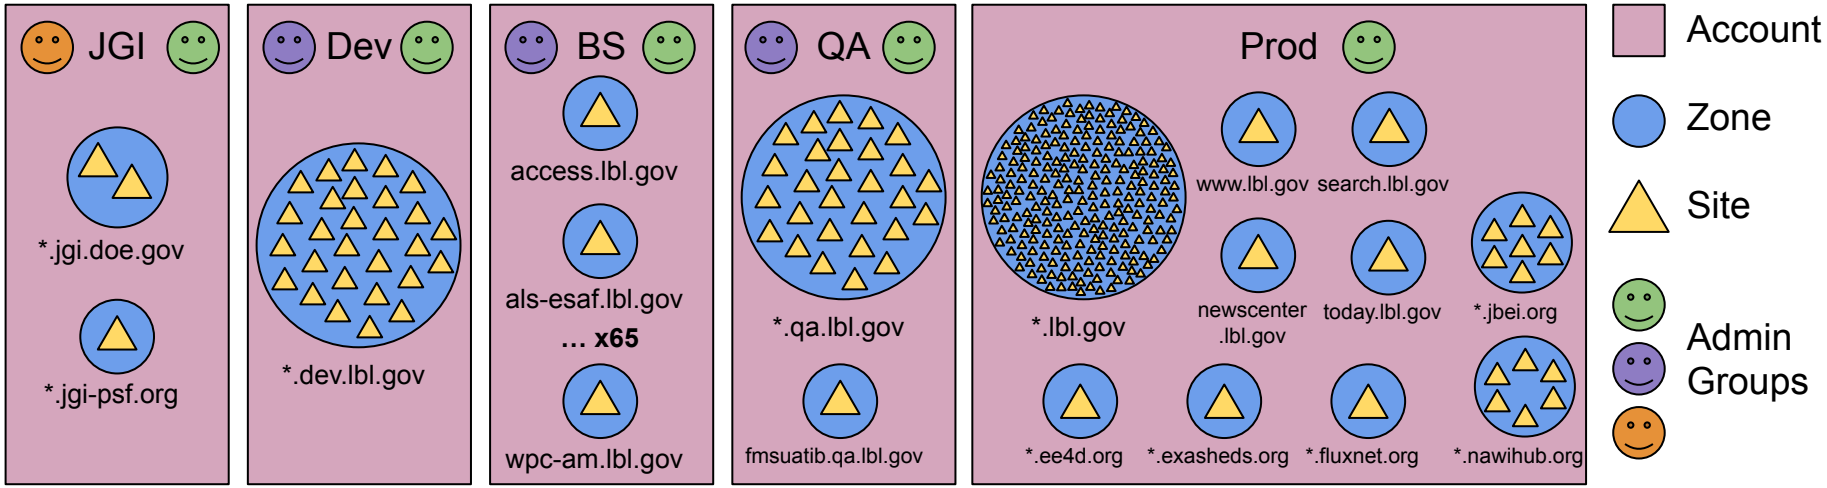

Known Limitations (4)

- Can only fetch 1 hour of logs from API at a time.
  - Must fetch per-Zone, can't just get all logs for your account with 1 query
  - Rate limit on API calls (it's changed over time)
  - Also size limit on response

- Some API endpoints require pagination (list of Zone IDs, for example).
  - Say we have 294 Zone IDs, can only fetch 50 "per-page", so have to loop over pages 6x to get all Zone IDs.

- Audit logs won't include login IP/timestamp for other users due to privacy reasons.

Lessons Learned (1)

- Gymnastics to accommodate non-web-only hosts…
  - "this server runs FTP, email, LDAP, etc etc in addition to our website"

- Issuing TLS certs for a new zone can be delayed, possibly 10-30 mins. If that downtime isn't acceptable, there's an API workaround to pre-provision the cert before go-live.

- Beware of CNAME chaining:
  example2 CNAME example1 CNAME example0
  If example0 moves to Cloudflare, must also create example1 and example2 in Cloudflare too or those will break.

Lessons Learned (2)

- Newly created Zones don't have log retention enabled by default (for API logpull). Can be enabled via API but not UI.

- Load Balancing default config may cause a large number of requests to the origin server. ~Mar 2019 we observed 2 requests from 164 different IPs every 60s.

Lessons Learned (3)

- Cloudflare uses CNAMEs (if you don't make them authoritative for DNS)
- CNAMEs can't coexist with any other record types, including SOAs (for apex records) or MX/TXTs (ex: email hosted at the domain/hostname)

✔️ example.lbl.gov CNAME example.lbl.gov.cdn.cloudflare.net

❌ example.lbl.gov CNAME example.lbl.gov.cdn.cloudflare.net
example.lbl.gov SOA nsx.lbl.gov. hostmaster.nsx.lbl.gov. 2019079980 14400 1800 2419200 300

❌ example.lbl.gov CNAME example.lbl.gov.cdn.cloudflare.net
example.lbl.gov MX 0 mailexchange.lbl.gov
example.lb.gov TXT "v=spf1 ip4:192.0.2.0/24 ip4:192.0.3.0/24 -all"

Lessons Learned (4)

- There is an unofficial workaround:

```
dig +short example.lbl.gov.cdn.cloudflare.net A
104.19.238.40
104.19.239.40
dig +short example.lbl.gov.cdn.cloudflare.net AAAA
2606:4700::6813:ee28
2606:4700::6813:ef28
```

❗ **These addresses could change!**
We haven't seen that happen since we started using CF (2yrs), but they could. We monitor for changes every 5 mins.

✔️
```
example.lbl.gov A 104.19.238.40
example.lbl.gov A 104.19.239.40
example.lbl.gov AAAA 2606:4700::6813:ee28
example.lbl.gov AAAA 2606:4700::6813:ef28
example.lbl.gov SOA nsx.lbl.gov. hostmaster.nsx.lbl.gov. 2019079980 14400 1800 2419200 300
example.lbl.gov MX      0 mailexchange.lbl.gov
example.lb.gov TXT      "v=spf1 ip4:192.0.2.0/24 ip4:192.0.3.0/24 -all"
```

Lessons Learned (5)

- If you're mass-migrating many sites, check to see they're actually functional before migrating to save yourself some headaches after ("why isn't this site working? Turns out it wasn't working beforehand, either")

- Some WAF rules are overly aggressive (too many FPs). May have to change to CAPTCHA challenge or disable entirely.

- WAF Rule ID 981176 is a generic ID used for the OWASP rule set indicating the threshold was exceeded. Need to drill down to see contributing rules.

1 Firewall > Overview > Activity log

Ray ID	5e565b40d9b1c7fe
Method	GET
HTTP Version	HTTP/1.1
Host	emp.lbl.gov
Path	/staff/andrew-satchwell
Query string	?page=4&s=year%2528%2529%2522%2526%25251%2527-%253B%253Cxss_62fdeda2b2556c46f9d6c093aa09805e%252F%253E%2527&o=desc&f%5Bauthor%5D=1447

2

Service	WAF
Rule ID	OWASP Block (981176) 3
Rule message	Inbound Anomaly Score Exceeded (Total Score: 31, SQLi=17, XSS=0) 4
Rule group	OWASP Inbound Blocking
Action taken	Log/Simulate

Export event JSON

5

Additional logs 7

950109 · Multiple URL Encoding Detected	OWASP Protocol Violations	Log	WAF
960024 · Meta-Character Anomaly Detection Alert - Repetative Non-Word Characters	OWASP Generic Attacks	Log	WAF
981318 · SQL Injection Attack: Common Injection Testing Detected	OWASP SQL Injection Attacks	Log	WAF
950001 · SQL Injection Attack	OWASP SQL Injection Attacks	Log	WAF
959073 · SQL Injection Attack	OWASP SQL Injection Attacks	Log	WAF
981245 · Detects basic SQL authentication bypass attempts 2/3	OWASP SQL Injection Attacks	Log	WAF
981243 · Detects classic SQL injection probings 2/2	OWASP SQL Injection Attacks	Log	WAF

Lessons Learned (6)

- Moving a site from one Zone to another requires a Cloudflare Support ticket and can take up to **two days** to complete.

- Changes may be needed for mail routing, including:
  - Updating MX records to point to the real Origin Server IP address, and not the hostname that resolves to Cloudflare (which will drop SMTP connections).
  - Updating mail routing rules (IronPorts, Google SMTP relay) to use the IP address or hostname that bypasses Cloudflare.

Lessons Learned (7)

- If you create a Zone and enable HSTS, any site moved there will be locked-in to HTTPS (up to the max-age, i.e. 1 year). Make sure HTTPS is working properly first in case you need to back-out of CF and still use HTTPS after.

- Client IP addresses for access control - requests will now come from CF IPs. CF can include the Client IP in HTTP requests, but Origin Server will need reconfig.

- Cloudflare adds complexity, some people will be hesitant to introduce that to their system. (Esp. with a couple recent worldwide outages)

Overall Experience

- Ease of Use
  - Shockingly easy sometimes
  - Very opaque at other times
- Performance
  - Only noticed caching weirdness (maps), generally unnoticeable
- Reliability
  - 2 major worldwide outages in recent memory
  - Very detailed post-mortems though
- Cloudflare Customer Service
  - Responsive, detailed technical answers
- Account Team
  - They do what they can, but we realize we're not a multi-million $ customer

Questions? Suggestions?

- mnsmitasin@lbl.gov
- cloudflare@lbl.gov

My API logpull script:
https://github.com/michaelsmitasin/lbl-cloudflare/blob/master/FETCH:Cloudflare-logs

Special Thanks to: Jay Krous, Phil Butler, Mark Dedlow