indeni

Check Point

Deployment Trends 2020

indeni

About this Report

Intended Audience

Current Check Point customers or those considering deploying Check Point Devices.

Publication Date

All information in this report is current as of the time of publication:

February 18, 2020.

Executive Summary

The purpose of this report is to provide a baseline of emerging and common practice in deployment and management of Check Point firewall products, based on characteristics and trends in production deployments in 2019.

Stakeholders in design and implementation of network security should consider the results of this report as guidelines for data-driven comparison against common industry practice. Specifically, this report will cover the following:

- What are the most commonly used versions of Check Point Software?
- What was the impact of End of Support for R77.30?
- What's "typical" across Security Infrastructure Operations teams?

Key Findings

R77.30 is (still) a commonly used Check Point software release.

R80.x has grown into the most commonly deployed version series, but R77.30 still is run in production at some sites, despite its transition to End of Support. Upgrades introduce new unknowns, and some organizations prefer to accept the known risk of lack of support.

R80.x shows quality improvements, at the cost of higher CPU utilization.

Compared to R77.x, R80.x shows many fewer issues with I/O and intra-cluster mismatches. However, there is a higher prevalence of high CPU utilization, and any issues with external identity services have a greater impact.

Customers are continuing to buy new Check Point hardware.

Key shifts in multiple metrics relating to firewall age show that, while many customers continue to use older hardware until it reaches hardware End of Support, there was a large influx of newer firewall hardware compared to last year's report.

Firewalls tend to experience fewer and/or less severe problems after OS upgrade.

Analysis of weighted sum of issues per day shows that Check Point firewalls tend to experience fewer and less severe outlying "bad days" after being upgraded from an older version to a newer version

What's "Average" for Security Infrastructure Operations?

The frequency and severity of issues among organizations shows a normal "bell curve" distribution on the stable low end, but on the chaotic high end there's a cluster of organizations forming their own flock of black swans.

Detailed Findings

Version Popularity

There are 18 release versions of Check Point running on devices in Indeni's study, with 87% of firewalls running either R77.30 or a release of R80.x. When the raw input is rationalized by dropping the outliers, that number rises to 95%. Unless otherwise stated, all percentages in this section represent the rationalized year-end market share of the respective version as of 2019-12-31, with raw numbers represented in parentheses (like this).

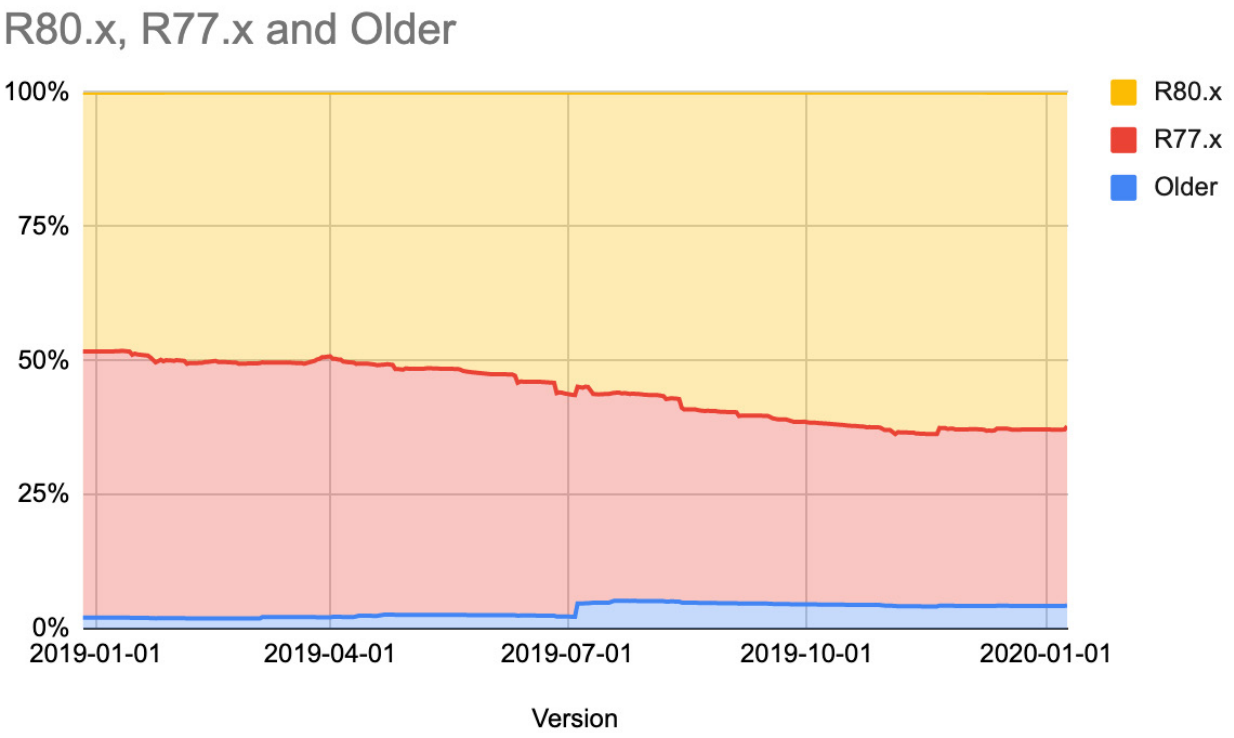

R77.30

At 33% (37%) of all firewalls monitored by Indeni, R77.30 continues to be a popular release in production, despite having reached EOS (End Of Support). However, this value is significantly down from its high of 65% in 2018, and is lower than the combined popularity of the R80 series at 63% (50%).

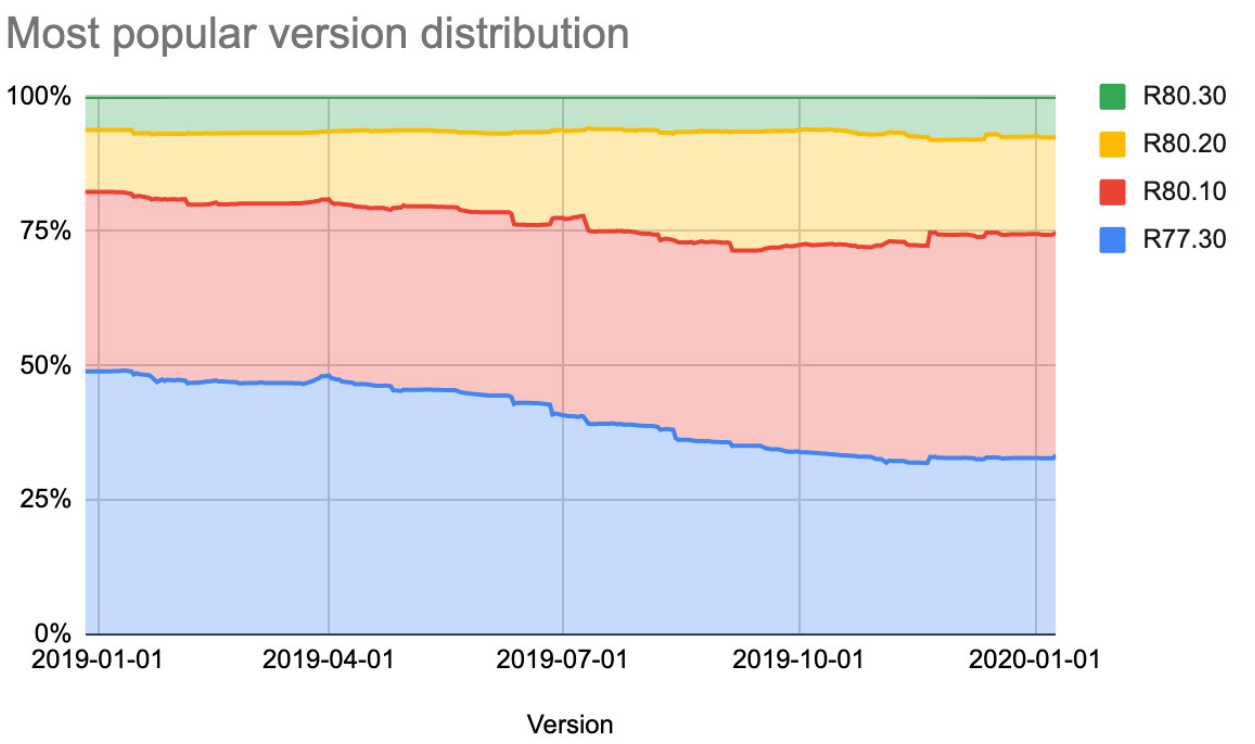

R80.x

In 2019, deployments of the R80.x series were predominant, representing 63% (50%) of all Check Point firewalls monitored by Indeni.

Last year, we predicted that the release of R80.20 would cap the market share growth of R80.10, and, while R80.20 did indeed experience significant growth, R80.10 more than doubled in market share to become the most popular deployment choice at 41% (28%).

R80.20 grew substantially to third place at 18% (17%), since it became Check Point's recommended version on January 15, 2019, although it has yet to eclipse R80.10.

R80.30 saw growth in 2019, although it remains relatively small at 7.5% (4.3%), and the release of R80.40 in January 2020 is likely to cap its adoption. Users who want an older, "proven" code base are likely to stick with R80.10, and R80.40 will likely attract users seeking "new features".

R80.40 was available via Early Availability in 2019, and showed up in trace amounts in the Indeni sample pool. Expect aggressive growth of this release in 2020 thanks to its dedicated HTTPS policy layer and zero-touch deployment.

Version Upgrades

The venerable R77.30 reached EOS (End of Support) in 2019, providing a strong incentive for customers to upgrade to the R80 series. Initially, EOS was scheduled for May 2019, with the date later deferred by Check Point to September 2019. While usage of R77.30 has decreased from last year, it is still run on a large number of firewalls that Indeni monitors. The version popularity distribution charts don't show a significant sustained decline in usage that corresponds with an increase in R80.x usage. In fact, 81.3% of the firewalls covered by this report ran the same version of Check Point for the entire year.

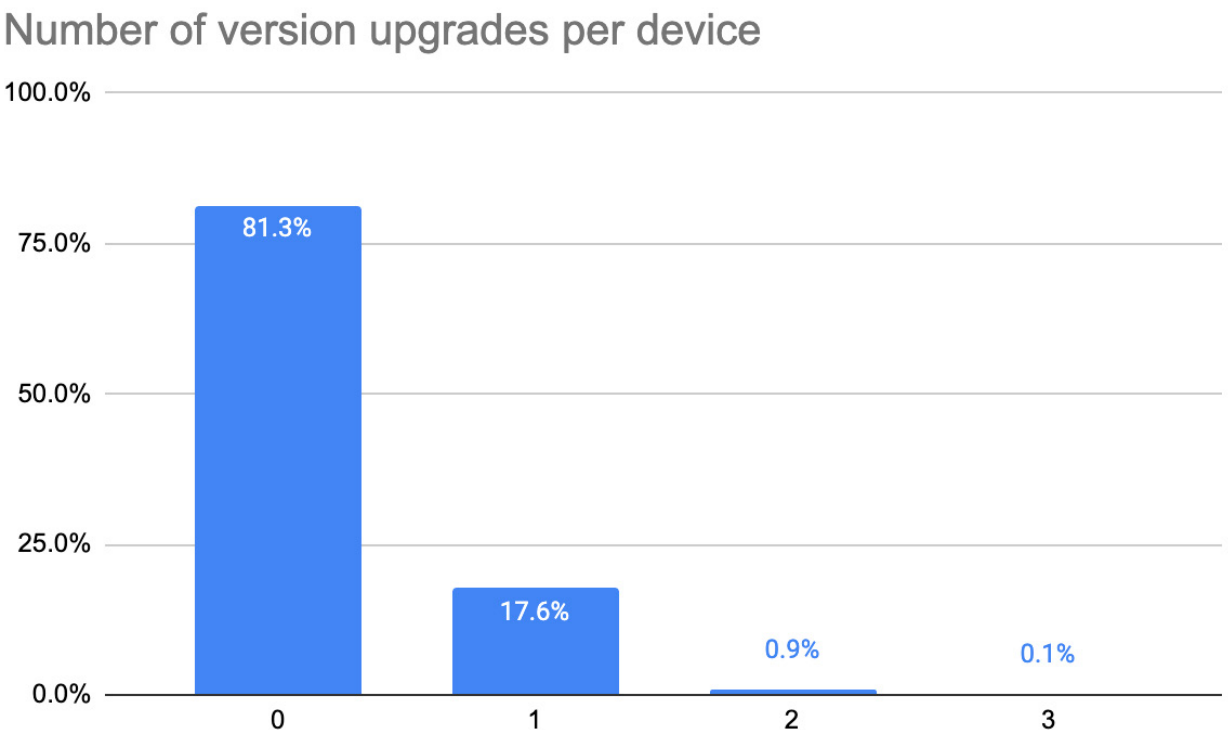

By focusing our results in this subsection on the subset of devices which were upgraded (~18%), certain trends emerge.

Upgrade Count From/To (2019)

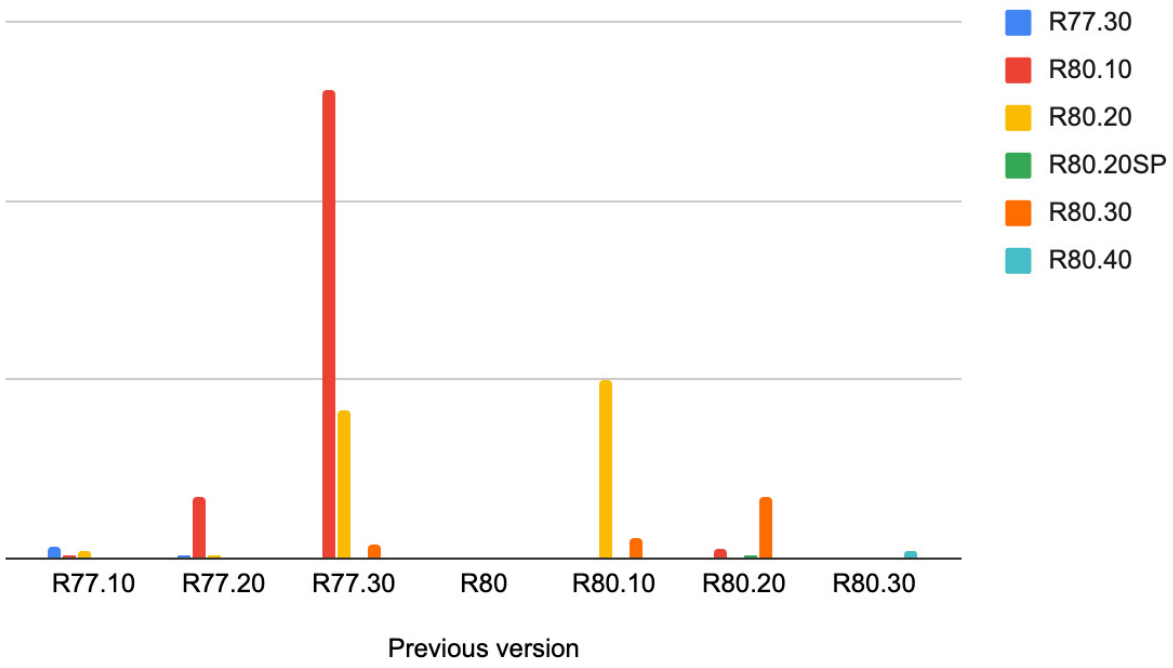

The largest migration pattern that Indeni observed during 2019 was from R77.30 to R80.10, with a distant second from R80.10 to R80.20, and a close third from R77.30 directly to R80.20.

Upgrades from R77.30 by date

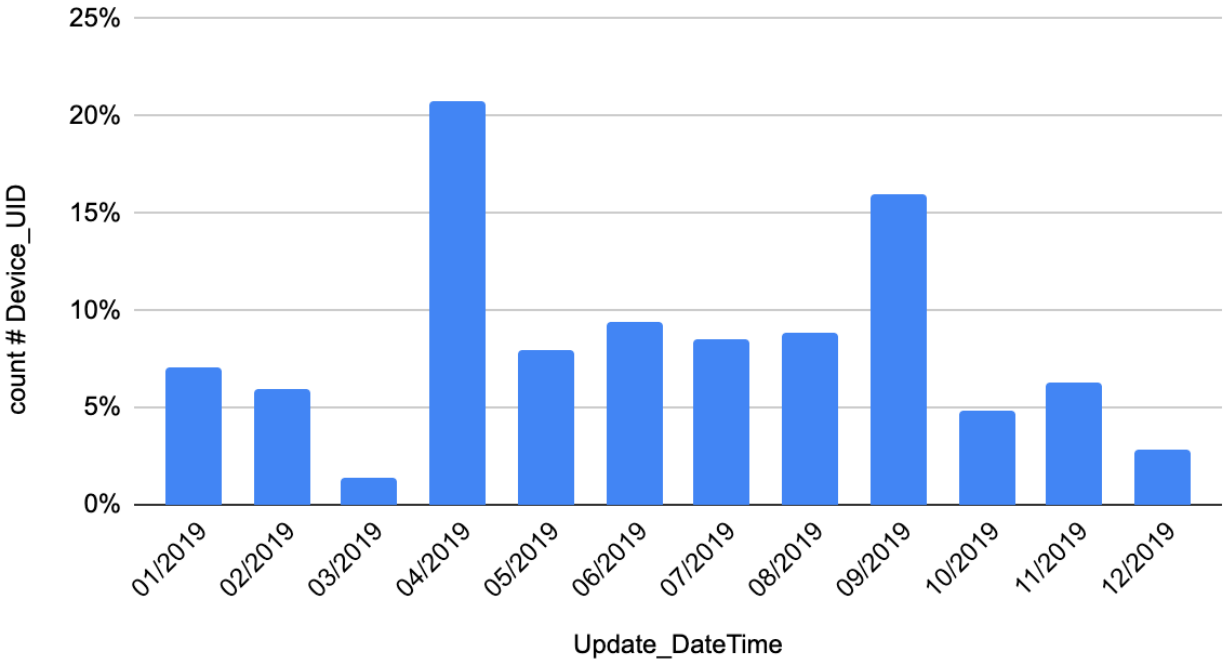

In terms of timing, the largest month for upgrades from R77.30 was 04/2019, the month before the first EOS date. The second largest month was 09/2019, the month of the deferred EOS. For the other months of the year, upgrades proceeded at nearly a uniform pace, indicating that, while the impact of EOS is measurable, it is not the only motivating factor in upgrading.

Indeni has previously studied the behavior of firewall upgrades[1]. After Palo Alto Networks released a security advisory and upgraded release for multiple major versions, Indeni found that there was no evidence that any devices we monitor had been upgraded 30 days after the announcement and new release. From that study, we determined that, since Indeni finds issues which are symptomatic of future outages, our customers tend to self-select into a profile of avoiding the unknown risks of upgrades by accepting current known risks[2]. In the case of R77.30, the known risk is losing access to support (unless extended support has explicitly been negotiated with Check Point).

The paradox of mission critical equipment is that, by definition, the organization depends on it. While many Information Security practitioners are familiar with the advice[3] to keep equipment at the latest patch level, patches are code changes which potentially introduce behavior changes, and behavior changes in mission critical infrastructure create widespread risk for the organization. Additionally, as firewalls gain deeper security inspection capabilities, they become more valuable sources of visibility and metrics, which enforces the version lock-in to avoid changes in metric generation method which would potentially invalidate point-in-time comparisons. From an organizational perspective, it is less disruptive to add compensating controls as new vulnerabilities are discovered. Once the upgrade becomes possible - or worse, necessary - all of the changes are absorbed at once, rather than numerous small changes experienced over time. If the impact of that One Big Change is highly disruptive, it often feeds back in a vicious cycle into fear of future changes, perpetuating the pattern.

1 https://indeni.com/blog/palo-alto-networks-pan-os-upgrades-after-a-vulnerability-announcement/

2 In preparing this report, Indeni learned that Check Point internally estimates R77.30 usage at less than half of Indeni's findings, adding evidence that Indeni's customer base may skew towards running older versions of equipment as a way of avoiding operational risk.

3 For examples see the NIST Cybersecurity Framework, the CIS Controls, the SANS Top 20, the Verizon DBIR, US-CERT, etc...

Top Issues by Version

While there is overlap in the issues found in monitored R77.x and R80.x firewalls, they tend to experience different "top" issues. R77.x experiences more interface-based issues like Tx/Rx throughput and Rx errors, as well as configuration mismatches between cluster members. R80.x experiences more high CPU usage and, notably, slow DNS response times.

With the increased focus on Zero Trust and HTTP inspection, it makes sense that R80.x would experience more issues of CPU usage and DNS lookup of FQDN-based objects. The lower incidence of configuration mismatches implies a quality improvement in cluster consistency, but the interface issues - normally a function of hardware - may be a result of newer devices which come pre-loaded with R80.x.

R77.x Issues	R80.x Issues
External Services • Next hop inaccessible • Certificate authority not accessible • NTP sync failure(s)	**External Services** • DNS server response time slow • BGP peer(s) down • Permanent VPN tunnel(s) down
Configuration/Administration • Jumbo Hotfix Take mismatch across cluster members • Critical configuration files mismatch across cluster members • DNS servers used do not match across cluster members	**Configuration/Administration** • Connected networks do not match across cluster members • Features enabled do not match across cluster members • Contract(s) expiration nearing
CPU/Resources • Aggressive Aging enabled • Interface nearing maximum Tx/Rx throughput • Rx packets experienced errors • Network port(s) down • Bond/LACP interface down	**CPU/Resources** • High CPU usage per core(s) • High load average • High CPU usage per core(s) • Some VSes have high CPU usage • High disk space utilization

Model Popularity

Over 50 different models or platforms were present in the Indeni Insight data. The most common appliances in this report are mid-range boxes in the 4000 and 5000 series.

New notable hardware in this year's study includes the MHO series (Maestro Hyperscale Orchestrators), although not yet in large numbers.

There were also significant numbers of VM instances across VMware, AWS, and OpenStack.

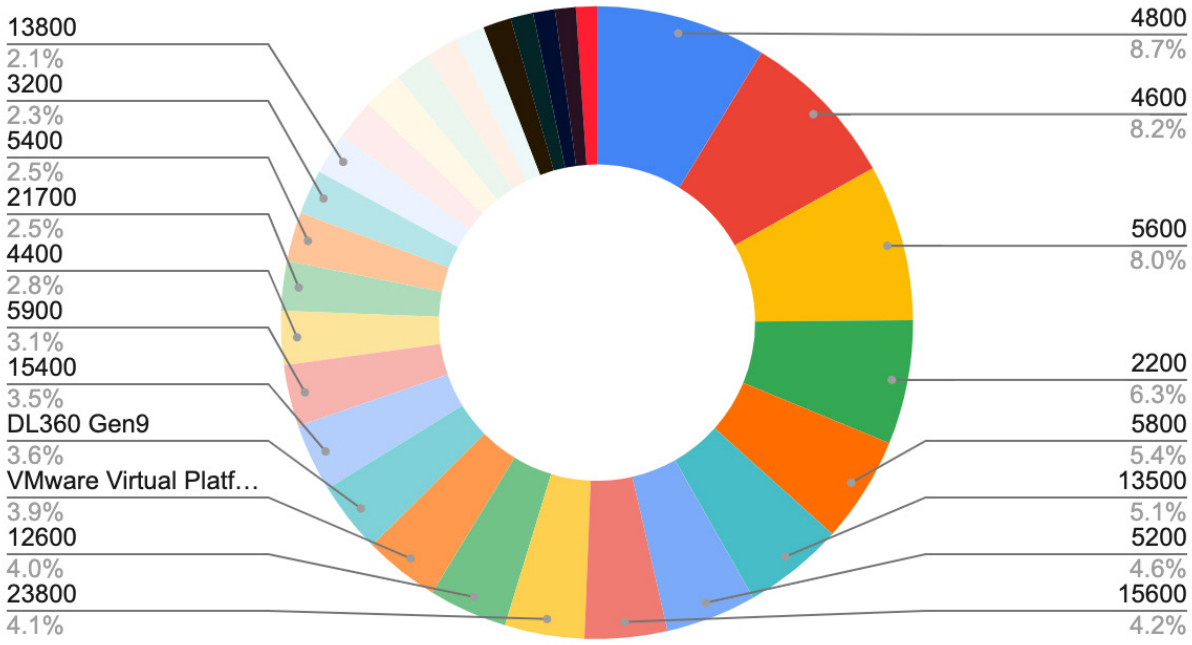

Devices by Market

Within the sample of Indeni customers, about 3/4 of the devices are marketed by Check Point[4] as Enterprise appliances, followed by 1/8 as High End DC (Data Center) appliances. The IAS-D series, about 4%, was a rebranding of HP ProLiant servers.

VMs account for 1/20 of total firewall instances. For this report, VMs are combined into the same Market category regardless of hypervisor, virtualization infrastructure, or cloud host. VMs are included as a convenient comparison point against firewall appliances and chassis, given that the hardware is abstracted.

In this report, the MHO series is included in the Scalable Platforms market, labeled here as “Scalable / Hyperscale”.

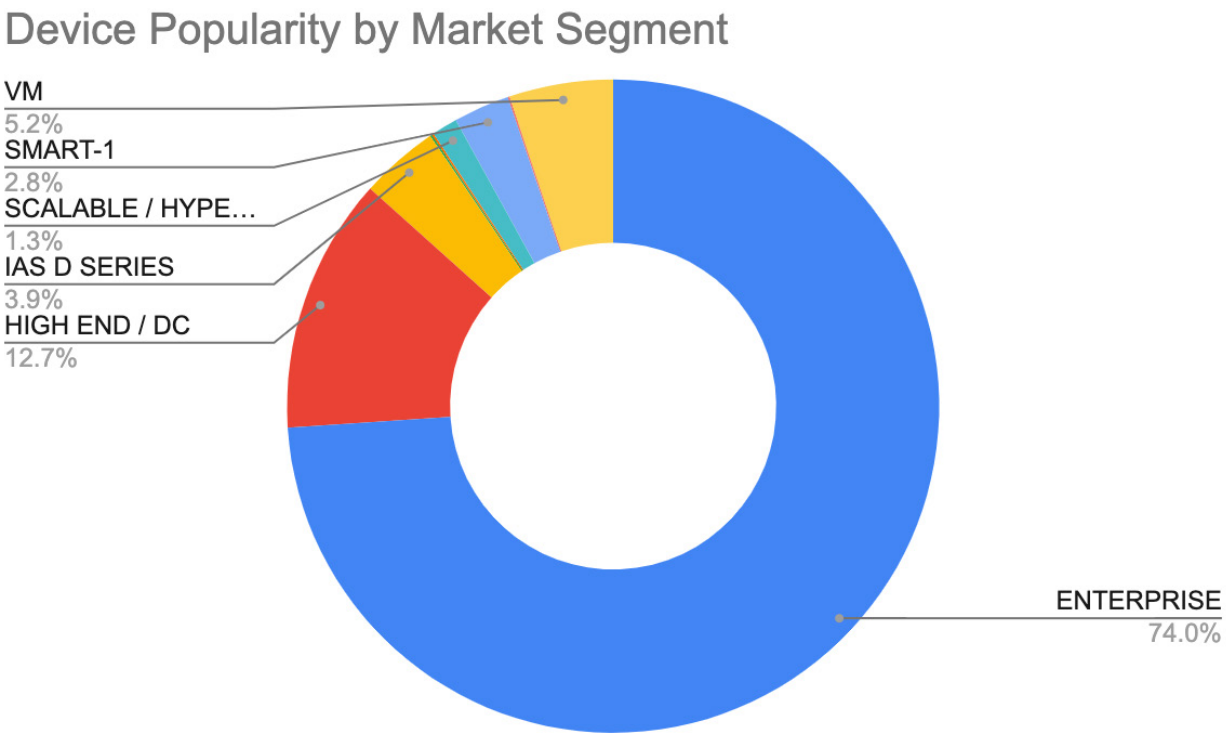

4 Classifications based on Check Point’s Support Life Cycle Policy web page as of 1/31/2020, https://www.checkpoint.com/support-services/support-life-cycle-policy/#AppliancesSupport

Devices by Age

There is a dichotomy among Check Point customers, based on the distribution of firewall ages: the majority are models released[5] either 3 years ago (27%) or 8 years ago (33%), These two data groups can be explained by looking at common IT equipment life cycle policies.

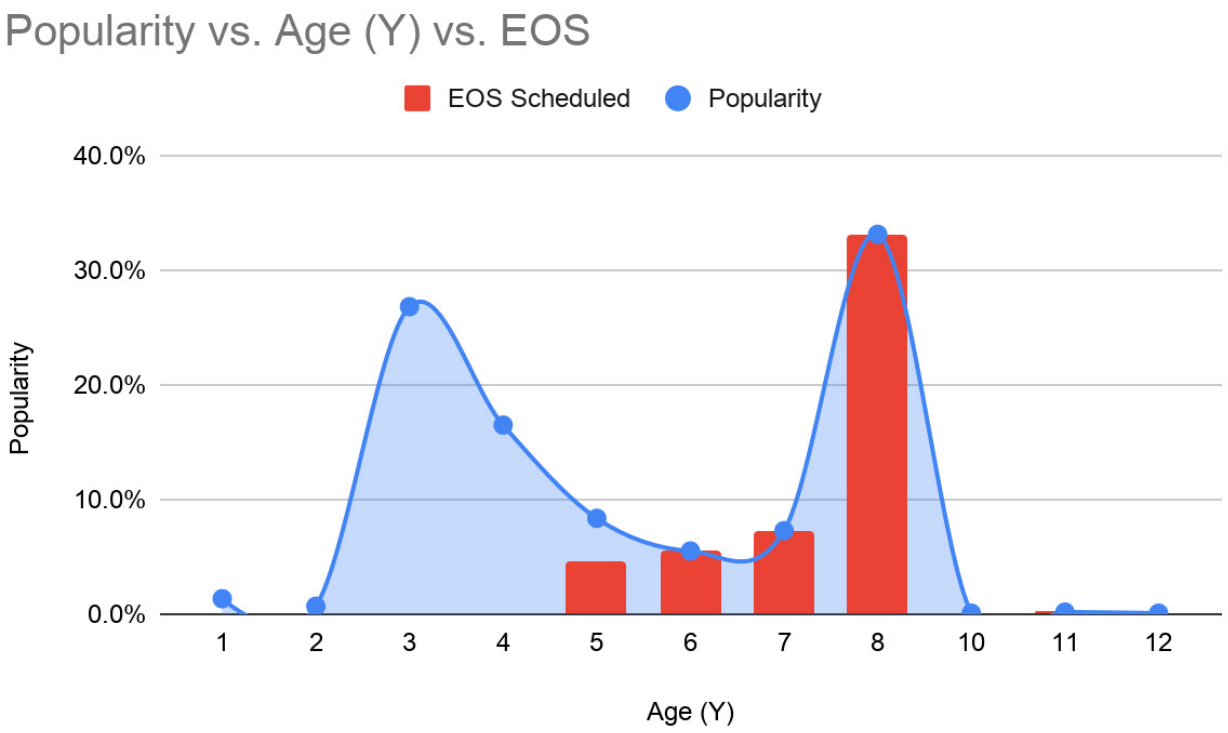

Production servers are often aged out after 3-5 years[6]. In this study, the group of firewalls at 3 years (27%) is followed by an exponential decline, which is consistent with those customers applying server lifecycle practices to their Check Point firewalls.

Network devices like switches are often expected to last 5-10 years or as long as the vendor provides support. In this study, the group of firewalls at 8 years (33%) immediately plummets, and the mere handful that are older are out of support. The long deployment lifetime and dearth of firewalls lacking support are consistent with network equipment lifecycle practices.

For purposes of comparison, we are including a Cumulative Age Distribution chart. This chart can be read from both top and bottom. For example, 46% of firewalls are at least 5 years old, and only 29% of firewalls are newer than 3 years old. Additionally, this chart demonstrates that the median device age is between 4 and 5 years old. This number represents a significant shift from last year's report, in which the median age was between 6 and 7 years.

---

5 Model release dates are defined as the "General Availability" date for each product as listed on the Check Point 'Support Life Cycle Policy' web page as of 1/31/2020: https://www.checkpoint.com/support-services/support-life-cycle-policy/#AppliancesSupport

6 Server life expectancy varies by industry and company, but publicly available data suggests the most common range is between 18 months and 5 years, with a median value of 3 years.

Cumulative Age Distribution

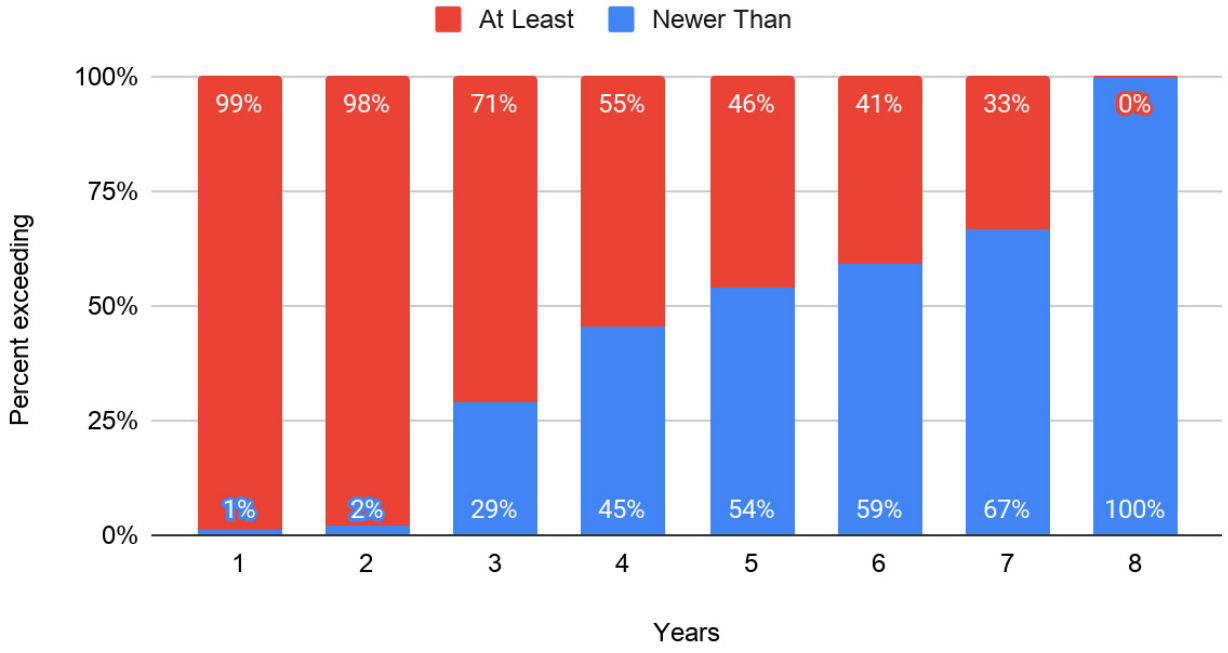

Examining age distribution by Market for the top two Markets show that “High End” devices follow a trend similar to the Enterprise devices, albeit skewing slightly older.

Device Popularity by Age and Market

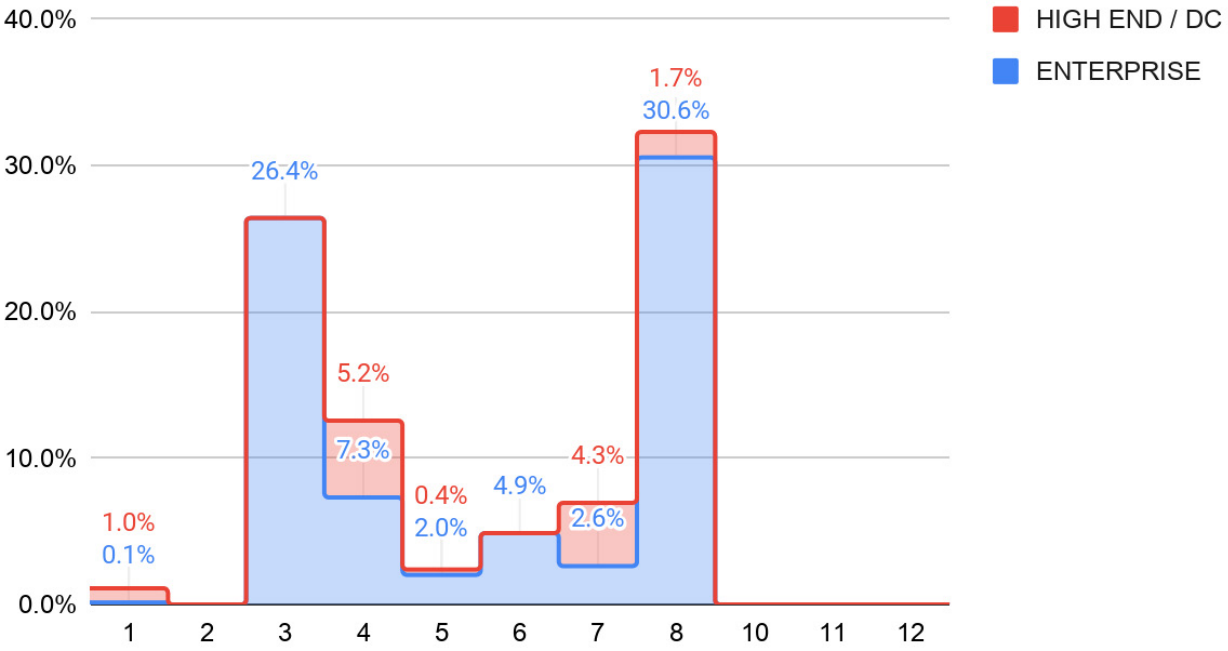

Devices with EOS (End of Support) Announcements

All models of Check Point appliances introduced in July 2014 or earlier have EOS dates scheduled. Approximately 45% of all firewalls in the Indeni study are scheduled for EOS in June 2022. This number is down from last year's report, at 52%.

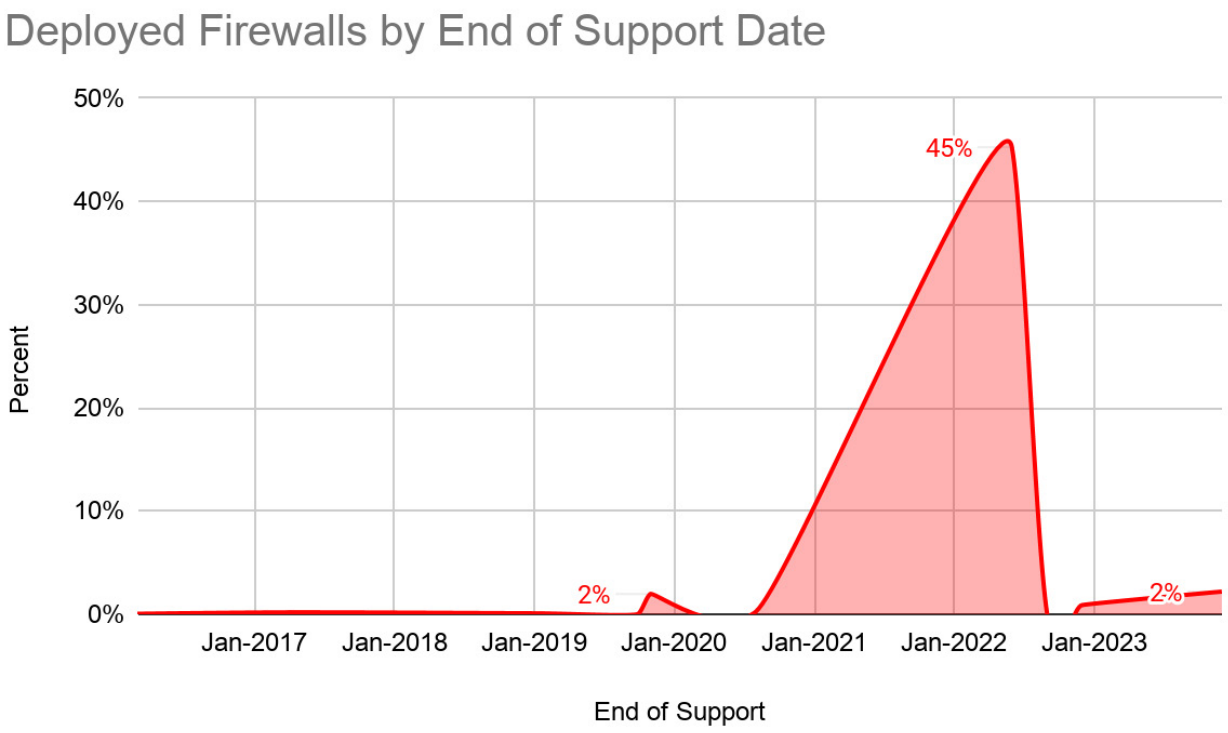

In comparing results of last year and this year, we do not expect to see a purchase crisis among Check Point users as June 2022 approaches. There are three factors leading to this conclusion. First, the number of devices affected by the June 2022 EOS date has decreased from last year. Second, although the second age peak has shifted from 7 years to 8 years, the younger age peak is still stable at 3 years, and has increased from 25% to 27%. Third, the median device age has shifted from 6-7 years down to 4-5 years. Taken together, there are indications of a strong pattern of purchasing newer equipment.This number represents a significant shift from last year's report, in which the median age was between 6 and 7 years.

Error Scores

This report uses an arbitrary “Error Score” to compare relative health and error rates between data groupings. The Error Score for a data grouping is the weighted sum of issues, normalized to remove bias from different member count between groups. While this methodology is by its nature reductive, stripping the context from each data point, it can be useful with a large enough sample size to produce a rough description of a “normal” expected experience for an arbitrary Check Point user, with a higher score representing more issues and/or more severe issues.

Indeni defines 4 levels of impact for issues:

1. CRITICAL is defined as “Device is down, or a current service disruption”.
2. ERROR means “You want to handle this within the next few hours”.
3. WARNING means “You’d want to clean this up as part of a project.”
4. INFO means “Just for your knowledge, up to you what to do with it.”

The weighting for Scores in this report is 1:2:4:8 from INFO to CRITICAL issue impact level.

Error Scores for Upgraded Devices

To examine the impact of a version upgrade, Indeni analyzed the telemetry from Check Point firewalls which were upgraded during 2019.

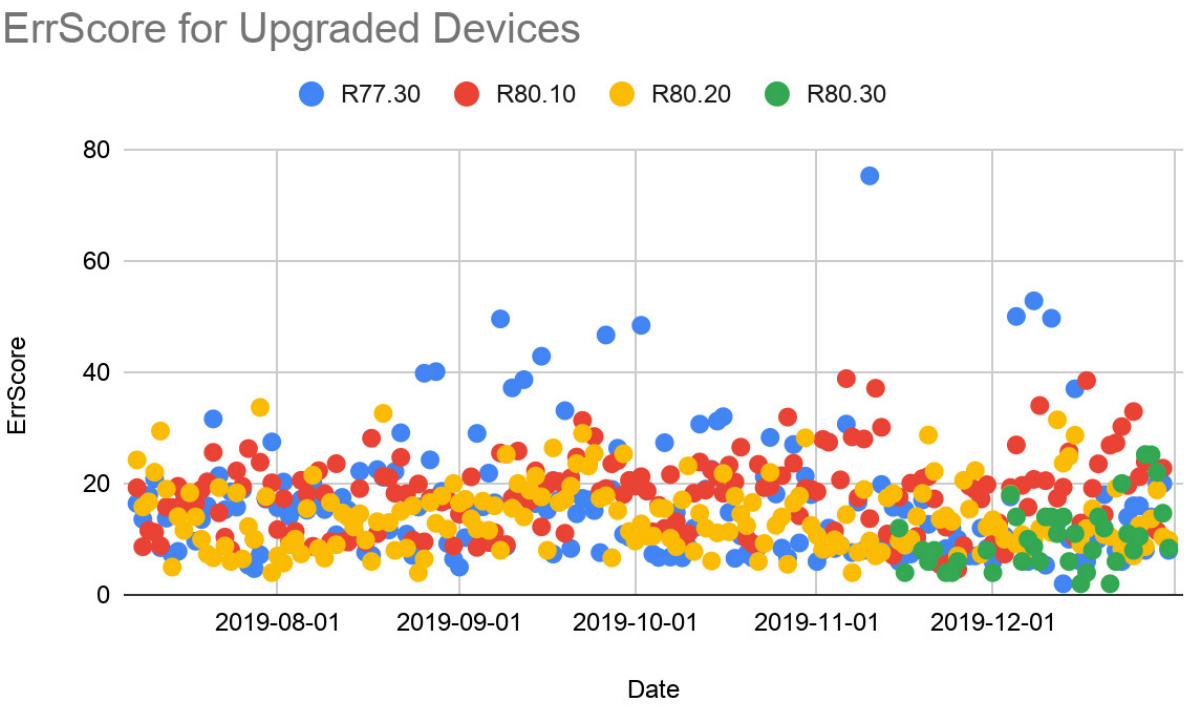

Results of this analysis show that the highest peaks occur for R77.30, with peaks decreasing with each subsequent release. Peaks represent outliers, the metaphorical "bad day at the office" during which some memorable event occurred. Therefore, lower peaks are likely to entail less disruption. However, the vast majority of data points fall within the same tight band, implying that there is little variance in the day-to-day operation of any of these versions.

Error Score Distribution across Multiple Organizations

There is a widely held belief in networking that no two networks are alike. Whether that's true, there's a question of how to measure the variance among networks, and whether they follow patterns that seem intuitive, like a classic "bell curve" Normal Distribution.

Indeni combined a normalized Error Score per Device per Day for each Indeni instance providing data for this report into a histogram to examine the distribution among different networks.

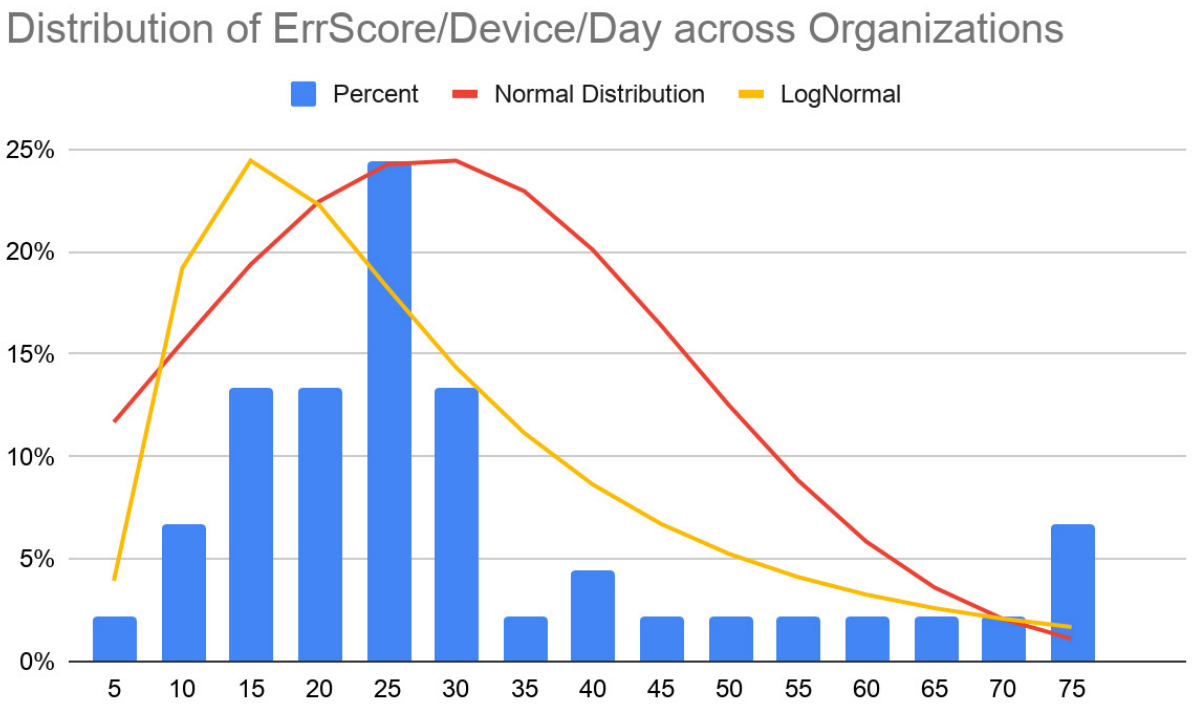

While the grouping of lower Error Score results follows a shape similar to a Bell Curve, there is a long tail among the higher scores, with an unexpectedly large grouping at the highest value. Unlike many data sets, like school exams scored between A-F or 0-100, Error Score is open ended: there is no theoretical maximum. Therefore, a long tail is not only possible, but potentially expected, despite its counterintuitive characteristics[7].

7 For a deeper dive into the counterintuitive characteristics of heavy tailed data sets, we recommend reading https://exploringpossibilityspace.blogspot.com/2013/08/tutorial-how-fat-tailed-probability.html

The implication on actual networks is that, from an operational stability perspective, many of them are indeed similar, but there's no network that's running so poorly it couldn't get worse. Even the best monitoring isn't enough to keep a network running well without the right knowledge and ongoing maintenance.

Example issues by category:

Best Practices Not in Place

- Hotfixes installed do not match
- SmartEvent log handling too slow
- Track CPU utilization
- Number of connections
- SNMPv2c/v1 used

Maintenance Tasks Pending

- BGP Peers down
- License Usage Nearing Limit
- Certificate nearing expiration
- Contract has expired
- Software support nearing expiration
- Hardware element down

High Availability Risks

- Static Routing Table does not match
- Network interface does not match across cluster members
- Bond Interface Down
- Management High-Availability Down
- Configuration changed by standby member

Network Visibility Challenges

- High CPU usage per core(s) for Check Point
- Static routing table does not match across cluster members
- Permanent/Monitored VPN tunnel(s) down
- Integration with identity/AAA server down
- High CPU Usage per Chassis and Blade
- Aggressive again enabled

Conclusion

Regardless of hardware or software in use, with proactive care, Check Point users can avoid many issues and reduce operational risk. Proactively implementing security, compliance, and vendor best practices early in the life cycle of Check Point devices will reduce tasks for IT operations and increase the return on investment in Check Point. By using Indeni, Check Point customers can proactively identify issues at setup, ensure the device-specific best practices are applied, and reduce the likelihood of future outages.

Getting started with Check Point Firewalls

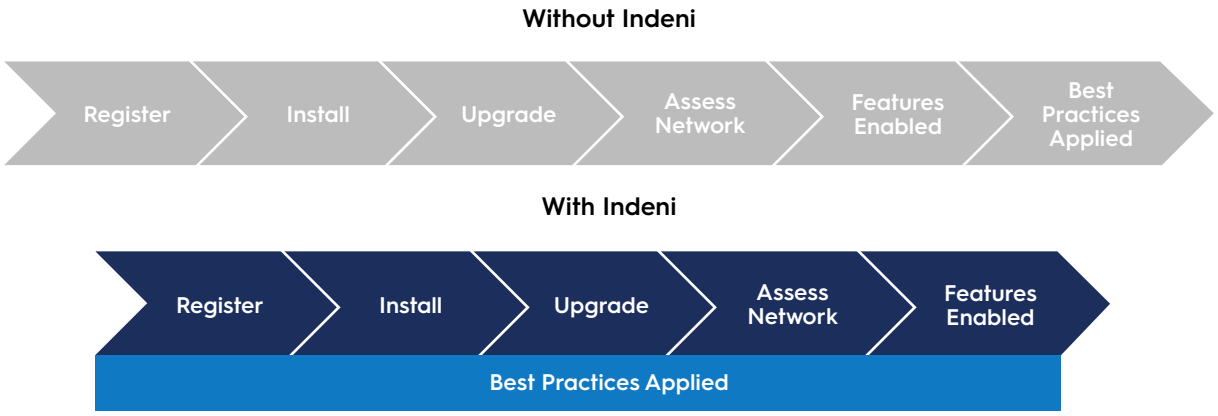

Join the Discussion

How do these findings align or differ from your environment? Ask questions, and give your feedback on this report in the Indeni Crowd Check Point Discussion forum.

Report Methodology

This report represents census data collected 1/1/2019 through 1/1/2020 from a pool of devices that were active 7/31/2019 or later at a subset of Check Point Customers who are also using Indeni Insight. All derived data were normalized per relevant grouping.

Indeni Insight analyzes billions of points of telemetry daily to offer a repository of device-specific intelligence focused on security infrastructure to aid in the development of turn-key operational automation and predictive analysis. This level of visibility across firewalls and other network and security devices allows Indeni to discover and provide remediation steps for issues that would otherwise go undetected. Learn more about the Indeni Automation Platform.

Corporate Headquarters
San Francisco, CA USA
Tel: +1-877-778-8991
Email: info@indeni.com

European Headquarters
Tel Aviv, Israel
Tel: +1-809-494-190
Email: info@indeni.com