

Nutanix Enterprise Cloud addresses the demanding requirements of Elasticsearch and other big data applications, reducing complexity, streamlining management, and improving TCO.

IHG Deploys Demanding Elasticsearch Application on Enterprise Cloud



InterContinental Hotels Group (IHG) is the second largest hotel company in the world with nearly 5,431 active properties in over 100 countries and an additional 1,776 hotels in the pipeline. For its U.S. operations, IHG has two main datacenters on the East and West coasts, supporting a variety of business-to-business (B2B) and business-to-consumer (B2C) applications, including important mobile apps.

IHG uses a combination of [Elasticsearch](#), Logstash, and Kibana ([Elastic Stack](#)) to provide real-time log analytics and metrics. The operations team relies on this data for operational visibility into applications and infrastructure; it helps them to ensure that everything is running as expected to meet availability SLAs. IHG's busy development teams also rely on data from Elasticsearch queries to troubleshoot code releases and user-reported bugs. Developers increasingly instrument their code to log data for these purposes, causing the amount of data ingested each day to grow at a rapid pace.

INFRASTRUCTURE CHALLENGES

Originally, Elasticsearch was deployed on a converged infrastructure solution. However, with the environment growing rapidly and becoming increasingly crucial to the business, a number of limitations became apparent:

- The converged solution was challenging and time-consuming to scale
- Software and firmware upgrades, maintenance, and other processes took too long and were too disruptive to users
- Total cost of ownership (TCO) was high

For these reasons, when it came time for a refresh the IHG team decided to look for a simpler, more cost-effective solution based on hyperconverged infrastructure (HCI).

CHOOSING NUTANIX ENTERPRISE CLOUD

The team evaluated several HCI solutions based on the following criteria:

- Performance
- Ease-of-use
- Management features
- Total cost

[Nutanix Enterprise Cloud](#) emerged as a clear leader. The industry's most popular HCI-based technology, Enterprise Cloud delivers a full infrastructure stack that integrates compute, virtualization, storage, networking and security to power any application, at any scale.

The Nutanix-native AHV hypervisor, included with each Nutanix purchase, enabled the IHG team to eliminate the VMware virtualization licensing costs that were a significant expense with the previous solution. Based on the Linux KVM hypervisor, Nutanix AHV provides full-featured virtualization that is simple to use because it is tightly integrated with the Nutanix Prism management plane.

“The lengthy maintenance outage windows were starting to impact our end users’ ability to access and utilize the valuable Elasticsearch data. Since the costs, time, and effort required for scaling had all reached an unacceptable level, we decided to look for a better solution.”

-Brian Malia, Systems and Storage Manager, IHG

“By using Nutanix Enterprise Cloud software as the platform for our Elasticsearch analytics engine, we are now able to collect a larger amount of data and analyze it in real-time. With better and faster analytics, we can be more proactive keeping our environment stable.”

-Brian Malia, Systems and Storage Manager, IHG

“The centralized management features are a definite plus with the Nutanix platform, as we can aggregate events, alerts, and performance monitoring for all of our sites into a single operations view. Plus, we no longer need to hire multiple IT subject matter experts, since anyone on our team can learn to manage Nutanix software in just a few hours.”

-Brian Malia, Systems and Storage Manager, IHG

DEPLOYING ELASTICSEARCH ON NUTANIX: BEST PRACTICES

To support its growing Elasticsearch solution, IHG deployed two identical six-node Nutanix clusters using a hybrid storage configuration combining SSDs and HDDs in each node. One cluster was deployed in each of IHG's two U.S. datacenters. Each cluster supports Elasticsearch logging for its datacenter. Users of the system can submit queries and search across all data if desired.

In addition to the migration from VMware ESXi to AHV, The IHG team also saw the move to Nutanix as an opportunity to upgrade to the latest versions of Linux and Elasticsearch. This was accomplished by running Nutanix in parallel with the original solution prior to cutover.

This approach also enabled the IHG team to work closely with Nutanix to optimize the Nutanix configuration for the actual production workload. Running in parallel was straightforward because IHG was already using Kafka to act as a message broker. Therefore, the Elasticsearch cluster running on Nutanix simply had to “subscribe” to the log data stream within [Kafka](#). Once subscribed, the old infrastructure and the Nutanix clusters received the same data in parallel.

Log data sets are kept online for a designated retention period before being aged off. The old infrastructure remained online until the Nutanix cluster had accumulated data for the full retention period, at which point the old infrastructure was retired.

During the optimization process, the IHG team working with Nutanix identified a number of best practices for their deployment of Elastic Stack running on Nutanix:

- Configure one Elasticsearch data node per physical node
- Stripe data across multiple vDisks
- Turn off Elastic shard replicas and Nutanix compression

There are a wide range of Elasticsearch use cases, so there is no single set of “one-size fits all” best practices. Anyone deploying Elasticsearch should carefully consider these best practices as part of their deployment plans, while also taking into account specific availability goals, I/O patterns, and any other unique requirements.

PLACING ELASTICSEARCH DATA NODES

Elasticsearch clusters rely on a number of [different node types](#). This segregation of functions is part of what makes Elasticsearch so scalable. In particular, Elasticsearch data nodes store data and perform all data-related operations including search.

For this deployment, the optimal configuration was to have a single data node on each physical node in the Nutanix cluster. This ensured that data nodes didn't compete with each other for the same I/O resources and maximized cache performance. Other node types can float within the Nutanix cluster without causing any ill effects.

Other Elasticsearch users have found that they can run 2-3 data nodes per Nutanix node, so the optimum is use-case dependent. A rack-id setting can be configured that helps with data placement and stops the co-location of data.

STRIPE DATA ACROSS MULTIPLE VDISHS

The IHG Elasticsearch environment is extremely write-intensive, with log data streaming in continuously from busy application servers. To achieve optimal I/O performance from hybrid storage, each Elasticsearch data node is provisioned with multiple sets of vDisks:

- Each vDisk is 200GB in size
- Data is striped across sets of five (5) vDisks
- Each data node VM currently has 7.5TB of storage

Additional capacity is always provisioned in sets of five vDisks as needed.

DISABLE ELASTIC SHARD REPLICAS AND NUTANIX COMPRESSION

Elasticsearch and Nutanix Enterprise Cloud use many of the same big data design principles such as replicas to protect data availability and compression. During tuning, it was noted that Elasticsearch was already compressing data. Therefore, Nutanix compression was disabled.

Similarly for replicas, if each solution is left in the default configuration it results in excessive redundancy and inefficient use of storage. Nutanix maintains a replication factor (RF) of 2 or 3 (tunable) for all data it stores. For this deployment, Elasticsearch shard replicas were turned off in favor of Nutanix replication.

NUTANIX ENTERPRISE CLOUD RESULTS

The IHG team saw a number of immediate benefits from moving its Elasticsearch cluster to Enterprise Cloud:

- **Real-time performance.** The Nutanix solution has allowed IHG to collect a larger amount of data and analyze it in real-time to improve operational stability and satisfy the growing demands of the user community.
- **Fast and easy upgrades.** With the converged infrastructure solution, even something as straightforward as a firmware upgrade took more than a week to accomplish. Nutanix 1-click upgrades for firmware, hypervisor software, and Nutanix operating software can accomplish the same task in a few hours. There's almost no impact on running applications and once started, there's nothing else administrators are required to do.
- **Improved infrastructure visibility.** Administrative teams can spend a significant percentage of their time just keeping tabs on infrastructure and performing routine tasks. Nutanix Prism Central allows the IHG team to monitor and manage everything in multiple Nutanix clusters, giving them back hours of time each week.
- **Simplified scaling.** With converged infrastructure solutions, scaling can be a challenge. You have to first identify the source of a bottleneck and then decide which resource is needed. Procuring and deploying the necessary resources and re-balancing loads consumes more time. With the hyperconverged Nutanix architecture, scaling is accomplished by adding additional nodes. Nutanix Prism can analyze compute, I/O and storage capacity trends and recommend the best configurations for new nodes, taking most of the guesswork out of the process.

The IHG team has been able to streamline its Elasticsearch operations in the face of unexpectedly rapid growth, achieving a substantial reduction in both capital expenses (CapEx) and operating expenses (OpEx).

“Moving our Elasticsearch deployment to Nutanix Acropolis and AHV has been an excellent decision for our company. We consider Nutanix to be a true partner to our organization—not just another technology provider.”

-Brian Malia, Systems and Storage Manager, IHG

LESSONS LEARNED AND FUTURE PLANS

At the time the IHG team undertook the move to Nutanix, they didn't fully appreciate how popular the Elasticsearch solution was becoming with the IHG developer and user community or how quickly it would grow. As with almost any infrastructure deployment, in hindsight there are a few things they would do differently:

- **Flash.** IHG's Elasticsearch workloads combine a heavy write workload with random reads. While the IHG team sees value in having an HDD tier for cold data that is aging out, they would definitely consider a larger flash tier or an all-flash configuration if they were starting over. Elastic itself [recommends that its indexes be written to flash storage for performance](#).
- **Memory.** In a similar vein, the team would configure physical nodes with more memory to support their intensive Elasticsearch use case. The current nodes have 256GB of memory per node.

One thing that IHG appreciates about the Nutanix architecture is the ability to combine nodes with different configurations in the same cluster. This gives them the opportunity to further optimize the infrastructure without having to rip and replace what they already have. They are currently in the process of adding some all-flash nodes to each of their Elasticsearch clusters to create a larger flash tier. They'll also be testing the impact of [AHV Turbo](#) (which was added in the Nutanix OS 5.5 release) which takes the I/O path out of the AHV kernel to achieve greater I/O parallelism. Using AHV Turbo Nutanix has been able to demonstrate the ability to deliver [1 million IOPS to a single VM](#).

For IHG, migrating its Elasticsearch deployment to Enterprise Cloud not only proved out the value of the Nutanix architecture it demonstrated the value of their entire Nutanix support team. The positive experience with Elasticsearch has resulted in the company moving more workloads to Nutanix. An additional Nutanix cluster was added in one datacenter during a hardware refresh. A variety of tier 2 and tier 3 applications have already been migrated to Nutanix with tier 1 applications still to come.

GETTING STARTED

To learn more about deploying Elasticsearch on Nutanix, read the solution note, [Elastic Stack on AHV](#). Whether you are looking to transform your approach to big data, or your entire IT operation, Nutanix serves as a trusted partner. You can try out Nutanix technology yourself by visiting nutanix.com/testdrive. To learn more about how Nutanix can help you, contact Nutanix at info@nutanix.com, follow us on Twitter [@nutanix](#), or send us a request at www.nutanix.com/demo to set up your own customized briefing.

“Nutanix is far less complex from a management perspective than the other solutions we tested. Before we moved to Nutanix, we had to coordinate all infrastructure changes with other IT groups in our company, and then upgrade every separate component one by one. Due to the need to involve multiple teams, infrastructure upgrades took over a week, compared to just a few hours with Nutanix.”

-Brian Malia, Systems and Storage Manager, IHG



T. 855.NUTANIX (855.688.2649) | F. 408.916.4039
info@nutanix.com | www.nutanix.com | [@nutanix](#)

Nutanix makes infrastructure invisible, elevating IT to focus on the applications and services that power their business. The Nutanix enterprise cloud platform leverages web-scale engineering and consumer-grade design to natively converge compute, virtualization and storage into a resilient, software-defined solution with rich machine intelligence. The result is predictable performance, cloud-like infrastructure consumption, robust security, and seamless application mobility for a broad range of enterprise applications. Learn more at www.nutanix.com or follow us on Twitter [@nutanix](#).

©2019 Nutanix, Inc. All rights reserved. Nutanix is a trademark of Nutanix, Inc., registered in the United States and other countries. All other brand names mentioned herein are for identification purposes only and may be the trademarks of their respective holder(s).