augustschell

CASE STUDY

U.S. Department of Defense Agency Juggles Disparate Data Centers and Aging Hardware

Executive Summary

A federal agency within the U.S. Department of Defense was tackling a three-year data center rebuild. In the midst of refreshing certificate systems and the network, it was discovered that they were experiencing repeated outages at one data center site with limited visibility into their IT program overall. At the same time, a hardware refresh was looming, while their Red Hat certificate system was at risk of breaking.

Challenges

- The customer was issuing customer requests from two geographically disparate data centers, but weren't able to serve requests from each automatically in the event that one went offline.
- Old hardware was presenting infrastructural limitations and a refresh was becoming urgent.
- Repeated downtime was causing frustration among customers.
- PKI credentials were being issued to over 5 million users. When authentication was attempted and the service wasn't available, authentication was required from the second site, but the customer couldn't access intelligence from both sites.
- Their three-tier application, a Red Hat certificate system, had a backend database. The Red Hat certificate system application was dependent on the database always being available, the customer had no insight into visibility of the application or database health.

Business Impacts

- ✓ The customer gained full visibility into the health of their applications.
- ✓ Downtime was reduced significantly, with each data center site able to clone its counterpart.
- ✓ The customer achieved service provider status, with the ability to remain online for all of their users and customers.
- ✓ A highly available, redundant data center design was put into place, resulting in a complete disaster recovery architecture.

Solutions

Primary

- BIG-IP Local Traffic Manager (LTM)

Supporting

- BIG-IP Application Security Manager (ASM)

DoD Agency Faces the Consequences of Potential Downtime and Aging Hardware

An agency of the DoD was running two separate, geographically disparate data centers, actively issuing customer requests from both sites. Due to the lack of integration between locations, the customer was unable to serve requests from one site in the event that the other went down.

"This customer was facing repeated downtime issues and they were receiving a lot of frustrated phone calls from their customers," Eric Hanlon, Senior Solutions Engineer at August Schell recalled.

Further, a complete hardware refresh was on the horizon, and their aging setup was presenting infrastructural limitations. "They were running on really old hardware," Hanlon explained. "They needed to do a total refresh and they were still using a copper based network infrastructure."

Preventing Downtime and Getting a Handle on Application Visibility

With two separate data centers to wrestle with, the customer wanted to be able to serve requests from the respective alternate site in the event that the primary site went down.

"They weren't aware of the scope of their applications, either," Hanlon said. "If they went down, no one would know, not even the ops team—and that was a huge problem. They'd have to go through the manual process of the customer calling in, reporting the outage, and then do an investigation to determine what the problem was."

With their ongoing downtime issues causing dissatisfaction among customers, plus the lack of intelligence between both data centers, the customer required a robust load balancing solution and better visibility into the status of their applications.

Ensuring the Health of Applications and Enabling Dual Site Replication

After identifying the primary needs of the agency, Hanlon and his team first set out to automate frequent application health checks. They wrote custom code designed to determine whether an application is up and "ask" if it's available every 15 seconds. If the application went down, it would automatically contact the other site for a 302 redirect.

"We spent a lot of time focused on integrating the F5 LTM product into their existing networking stacks. It became a full proxy sitting in the middle of the traffic flow with custom health monitors that execute specific flows to the Red Hat certificate system. The goal is to provide a known answer on health—this way, you can ask the app if a component is there, and it'll tell you the known response we hardcoded into HCS to make it app-specific," Hanlon explained.

Once the health monitors were up and running for the front end, the team replicated the database between sites and set up a secure VPN between them through F5, using the IP SEC tunnel for database replication. When changes were made at one site, they would seamlessly replicate to the other site, making each site aware of app-specific information.

"We used the same network to also do app monitoring and make intelligent routing decisions," Hanlon said.

Finally, the customer added the BIP-IP Application Security Manager (ASM) to the environment. "We made that addition for layer 7 payload inspections. The ASM sat on the F5, so if there was an issue with Cisco, for example, they still had F5. It allowed a security model where they could block traffic that clearly was malicious traffic they whitelisted through the perimeter firewalls," Hanlon said.

F5 Brings Depth of Network Traffic Intelligence, Protects Apps and Data

With the project complete, the customer was relieved to gain visibility into their applications from the F5 at the networking level.

"They weren't operating blind anymore. They were happy, and it made their lives a lot easier," Hanlon concluded.

The customer was set up to be notified if an application was down, and it would be handled by F5 in real time, giving them time to troubleshoot and reduce the potential for unmanageable consequences.

Looking forward, the August Schell engineering team plans to assist the customer with TLS break and inspect in order to separate malicious traffic and take action on it. In the future, the customer will be enabled to review protocol system traffic from user web browsers, VPN traffic, etc., and separate out malware, botnets, and file transfer.

51 Monroe Street, Suite 1802
Rockville, MD 20850
301.838.9470 | www.augustschell.com