BigFix: Scaling Relays Up, and Bringing Costs Down

Authors: Mark Leitch, Massimo Marra, Davide Cosentino @ HCL Technologies

BigFix is a powerful and precise security product capable of managing hundreds of thousands of endpoints. The BigFix relay infrastructure provides a superior, scalable, hierarchical model for effectively managing hundreds of thousands of endpoints. In terms of pure capacity planning, the historical model would have each BigFix leaf relays managing up to 1,000 endpoints.

With the release of BigFix 9.5.9, the Windows and Linux based leaf relays can now manage up to 5,000 endpoints. The impact of this 500% scalability improvement is immediately obvious: BigFix can be deployed with an 80% reduction in leaf node infrastructure requirements. This means reduced cost to deploy and manage, ultimately improving the exemplary ease of deployment that BigFix is known for.

We will provide a deeper overview of BigFix relays, and why they are more effective than ever in managing massive, worldwide enterprise deployments.

BigFix Relays: Digging Deeper

BigFix relays offer a broad spectrum of function and configuration options. Some example of the BigFix relay capability follow.

- Dynamic routing through relay affiliation and failover mechanisms.
- Adaptive bandwidth throttling.
- Proxy management.
- Cache management.
- Peer nest management (aka PeerNest).

These features go well beyond the theoretical: BigFix has thousands of deployments and the relay infrastructure has been able to adapt and thrive in diverse, real world production environments. The core mathematical model of the relay infrastructure is the k-ary tree: essentially a tree structure where each node in the tree can have up to “k” children. While the value of “k” will vary by deployment, this is an extremely effective and efficient mechanism. These trees have some interesting properties. For example:

- Most people are familiar with binary trees. A binary tree is simply an instance of a k-ary tree where k=2.
- The height of the tree determines the number of hops needed to traverse to a leaf node.

- The formula for the height of a perfect k-ary tree, where "N" is the total number of nodes, is as follows:

$$h = \lceil \log_k((k-1) \cdot N + 1) - 1 \rceil$$

Why is this so powerful? For example, for a 7-ary tree comprised of 49 leaf nodes supporting 245,000 endpoints, the tree height is merely "2". This essentially means it is only three "hops" to traverse over the entire enterprise! This is why tree implementations are used wherever rapid scale is required. This ability to manage rapidly and at high scale is accelerated via the BigFix relay scalability improvements! This benefits not only the agent management and content distribution capability of BigFix, but also the BigFix Query functionality as well!

BigFix Relays: Deployment Consideration

With the BigFix relay scalability improvements there are some deployment considerations.

- Relays virtualize extremely well. Virtual deployments are not only possible, but encouraged.
- The relay infrastructure must still be managed. Network limits, operating system configuration, etc. will influence deployment capability.
- A general best practice is to implement a "false" root. The principal is for a dedicated relay to "impersonate" the root server, to offload traffic and ensure health of the root server.

The following figure shows a sample 5-ary tree deployment with a false root implementation. The references at the end of this blog provide comprehensive information for managing relays.

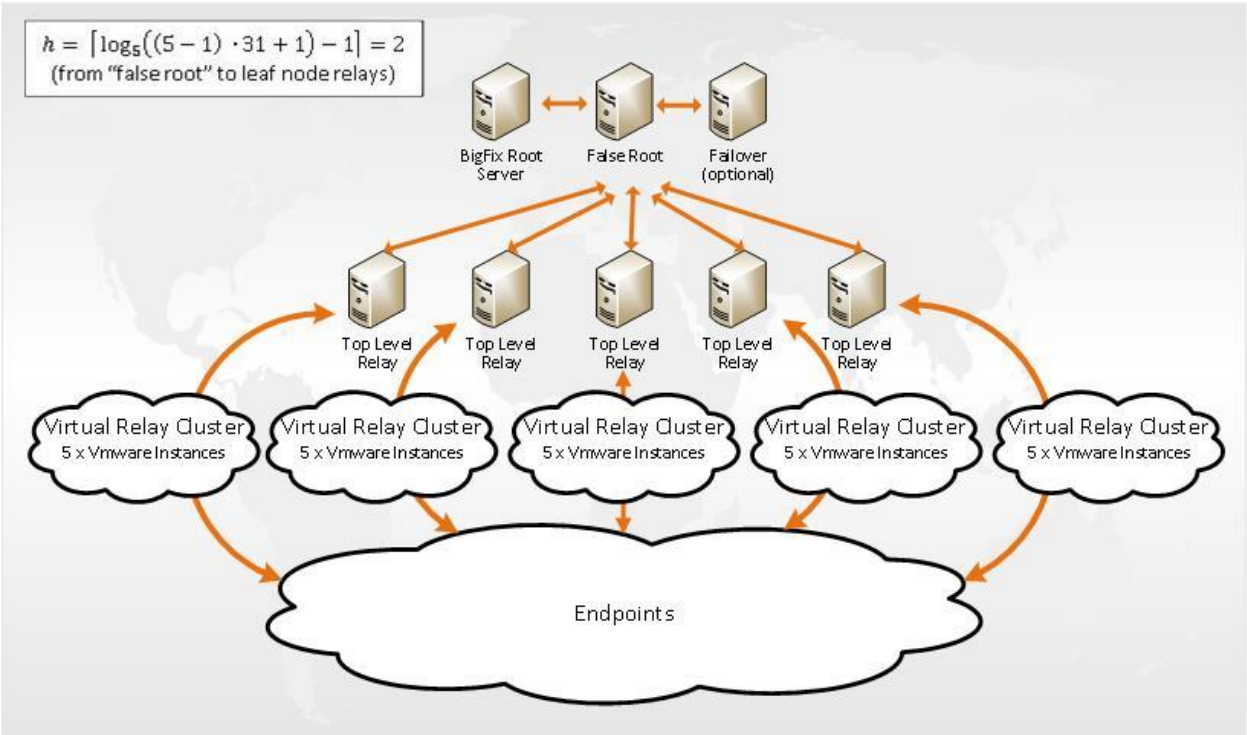

Further Reading

In the event further reading is desired on BigFix relays, and managing BigFix at scale, the following technical resources are available.

BigFix Platform Documentation: URL
BigFix Configuration Settings: URL
BigFix PeerNest: URL
BigFix Capacity Planning Guide, Maintenance Guide, and Performance Toolkit: URL