BFF: Complexity and Operational Risk

Lori MacVittie, 2011-14-12

#adcfw The reason bars place bouncers at the door is because it's easier and less riskier to prevent entry than to root out later

No one ever said choosing a career in IT was going to be easy, but no one said it had to be so hard you'd be banging your head on the desk, either. One of the reasons IT practitioners end up with large, red welts on their foreheads is because data centers tend to become more, not less, complex and along with complexity comes operational risk. Security, performance, availability. These three inseparable issues often stem not from vulnerabilities or poorly written applications but merely from the complexity of data center network architectures needed to support the varying needs of both the business and IT.

Unfortunately it is often the case that as emerging technologies creep (and sometimes run headfirst) into the data center the network is overlooked as a potential source of risk in supporting these new technologies. Traditionally, network readiness has entailed some load testing to ensure adequate bandwidth to support a new application, but rarely is it the case that we take a look at the actual architecture of the network and its services to determine if it is able to support new applications and initiatives. It's the old "this is the way we do this" mantra that often ends up being the source of operational failure.

COMPLEXITY MEANS MULTIPLE POINTS of FAILURE

Consider the simple case of using SPAN ports to mirror traffic, a traditional network architecture technique that attempts to support the need for visibility into network traffic for security purposes without impeding performance. SPAN ports are used to clone all traffic, allowing it to traverse its intended path to an application service while simultaneously being examined for malicious and/or anomalous content. This architectural approach can inadvertently cause operational failure under heavy load – whether caused by an attack or a flash-mob of legitimate users.

> "One of the problems with SPAN ports, which people tend to use because they're cheap, is that you won't get to keep it for your use all the time. Someone will come along and need that because there's a limited number of them," said John Kindervag, senior analyst with Forrester Research. "Whenever you are under attack and need that data, the switch is going to get saturated and the first port that quits functioning is the SPAN port so that it can have some extra compute capacity. So at the exact time that you need it, the whole system is designed not to get that data to you."
>
> …
>
> Network traffic capture systems can perform media conversion, sending lower bandwidth data streams to these network security appliances. They can also load balance these data streams across multiple appliances.
>
> -- Network traffic capture systems offer broader security visibility

Unfortunately, it isn't only SPAN ports (and thus the systems that rely upon them) that fail under load. Firewalls, too, have consistently failed under the arduous network conditions that occur during an attack. The failure of these components is devastating, disruptive, and unacceptable and is caused in part by architectural complexity. It isn't, after all, the IPS or IDS that's failing – it's the port on the switch upon which it relies for data. The applications have not failed, but if the firewall melts down it really doesn't matter – to the end user, that's failure.

One of the ways in which we can redress this operational risk is to simplify – to reduce the number of potential points of failure and more intelligently route traffic through the network.

INSPECT at the EDGE

If you look at the cause of failures in network architectures there are two distinct sources: connections and traffic. The latter is simple and it is also well understood. Too much traffic can overwhelms network components (and you can bet if it's overwhelming a network component, it can easily overwhelm an application), introducing errors, high latency, lost packets, and more. This trickles up to the application, potentially causing time-outs, unacceptably long response times, or worse – a crash. The answer to this problem is either (1) increase switching capacity or (2) decrease traffic.

Both are valid approaches. The latter, however, is rarely the path taken because it's it's an accepted fact that traffic is going to increase over time and there's only so many optimizations you can make in the network architecture that will optimize it and decrease traffic on the wire.

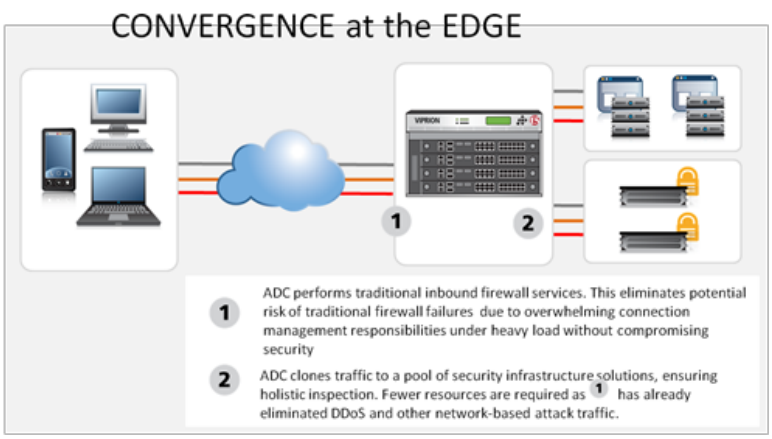

But there are ***architectural*** changes that can decrease traffic on the wire. It makes very little sense to expend compute and network processing power on traffic that is malicious in nature. The risk to application servers and network infrastructure is real, but avoidable. You need to check the traffic ***at the door*** and only let valid traffic in, otherwise you're going to end up expending a lot more resources tracking it down and figuring out how to kick it out. The problem is that the current network bouncer, the firewall, isn't able to adequately detect the fake IDs presented by application layer attacks. This traffic just slips through and ends up causing problems in the application tier.

The closer to the edge of the network you are able to detect – and subsequently reject – malicious traffic the more operational risk you can mitigate. The more assurances you can provide that security infrastructure won't end up being ineffective due to a SPAN port shut down, the more operational risk you can mitigate. Leveraging a converged approach will provide that assurance, as well as the ability to sniff out at the door those fake IDs presented by application layer attacks.

Leveraging a network component that is capable of both detecting inbound attacks as well as cloning traffic to ensure holistic inspection by security infrastructure will reduce complexity in the network architecture and improve the overall security posture. Detecting attacks at the very edge of the network – network ***and*** application-layer attacks – means less burden on supporting security network infrastructure (like switches with SPAN ports) because less traffic is getting through the door. If the network component is also designed to manage connections at high-scale, then the risk of firewall failure from overwhelming inbound connections that appear legitimate but are not.

Emerging architectural models are based on the premise of leveraging strategic points of control within the network; those places where traffic and flows are naturally aggregated and disaggregated through the use of network virtualization. Leveraging these points of control are critical to ensuring the success of new architectural and operational deployment models in the data center that allow organizations to realize their benefits of cost savings and operational efficiency. The application delivery tier is a strategic point of control in the new data center paradigm. It affords organizations a flexible, scalable tier of control that can efficiently address all three components of operational risk.

Consolidating inbound security with application delivery at the edge of the network makes good operational sense. It reduces operational risk across a variety of components by eliminating the complexity in the underlying architecture. Simplification leads to fewer points of failure because complexity and operational risk really are BFF – you can't address one without addressing the other.

- F5 Friday: The Art of Efficient Defense
- F5 Friday: When Firewalls Fail…
- F5 Friday: Performance, Throughput and DPS
- The Pythagorean Theorem of Operational Risk
- At the Intersection of Cloud and Control…
- When the Data Center is Under Siege Don't Forget to Watch Under the Floor
- Challenging the Firewall Data Center Dogma
- What CIOs Can Learn from the Spartans

- What is a Strategic Point of Control Anyway?
- Server Virtualization versus Server Virtualization

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com