Hillstone NETWORKS

WHITE PAPER

# Hillstone Twin-Mode Firewall Solution for **Redundant Data Centers**

# Introduction

As more enterprises rely on 24/7/365 application availability, redundant data center design with application failover has become increasingly widespread. Until recently, however, even the most sophisticated firewalls were incapable of ensuring full security in a redundant data center environment. Hillstone's Next Generation Firewalls address this issue head-on with the Firewall Twin-Mode feature, which links redundant firewall pairs across data centers to maintain full security for all redundant data center traffic flows. With Hillstone Firewall Twin-Mode, enterprises can achieve 24/7/365 business continuity and disaster recovery without any gaps in data center security.

## Contents

**Learn more about the Hillstone Networks Comprehensive Infrastructure Protection solutions**

Visit us at **hillstonenet.com**

© 2021 Hillstone Networks All Reght Reserved

# Redundant Data Centers

Redundant data center design has become increasingly mission critical as enterprises depend on applications and the network to fulfill key business functions. Industry sectors such as financial services, healthcare, and service providers depend on 24/7/365 application uptime, performance and security to function, as any loss of application access or data can seriously damage revenue and reputation. Especially in the case of healthcare, downtime can cost human lives.

**A redundant, or active/active, data center configuration is inherently more efficient than legacy active/passive mode operation, in which systems in one data center sit idle until there is a need for system failover.**

Redundant data center design mitigates the impact of hardware, software and data center failures so that systems and the business can continue to operate and generate revenue on a 24/7/365 basis. It incorporates a number of technologies, including server virtualization and application load balancing. Many firewall solutions incorporate clustering features that failover security services to a second active data center firewall if the first one goes down for any reason. However, natural disasters, terrorist attacks, or human error can bring down data center applications or an entire data center, requiring application failover to a second geographically dispersed data center to maintain uptime.

With a redundant data center architecture, the same critical enterprise business systems run in at least two data center sites simultaneously, with both data centers serving users continuously and functioning as backups for one another. When an application in one data center fails, the second or third data center takes over, providing services to all business users without interruption. The application switchover is instant and completely invisible to the user, who simply continues working as usual.

A redundant, or active/active, data center configuration is inherently more efficient than legacy active/passive mode operation, in which systems in one data center sit idle until there is a need for system failover. Redundant configuration doubles the capacity of a single data center through resource integration.

© 2021 Hillstone Networks All Reght Reserved

# Stateful Firewall Failover Issues

A redundant data center architecture typically harnesses special Data Center Interconnection (DCI) devices to extend the internal LAN across data center sites. This allows virtual servers to migrate across them effortlessly for disaster recovery, business continuity and resource allocation.

Most data center hardware devices, such as routers and load balancers, support redundant data center failover with one crucial exception: stateful firewalls. The reason? Stateful firewalls need to analyze the state of all session information to apply security policy effectively.

If one part of a stateful firewall session attempts to traverse a different firewall than the one where it originated, the second firewall will be unaware of the session established on the first. The second firewall will most likely drop the data flow, killing the transaction. This situation is known as asymmetric traffic flow, and may occur in the following scenarios:

- A **distributed services deployment** in which a few applications or services run in only one of the two redundant data centers, rather than in both concurrently.
- A **backup scenario** in which a second data center has taken over a business application that failed in the first data center.
- A **virtual machine migration** across the extended LAN from one data center to another. Each data center is running a pair of active firewalls to handle data security and serve as a standby for each other in case one fails.
- A transaction sends traffic from Service 1 servers in data center A (on the left) to Service 3 servers in data center B (right).
- IT policy requires all Data Center A traffic to traverse Data Center A firewalls (see green line). Data Center A firewalls establish a stateful inspection session.
- The traffic travels across the DCI LAN extension devices to Data Center B.
- When the requested data is returned by servers in Data Center B, IT policy requires it to traverse Data Center B firewalls.
- Unfortunately Data Center B firewalls are unaware of the session established on Data Center A firewalls. By default they drop the return data, killing the transaction (see red line). The requested information is never delivered back to the user.

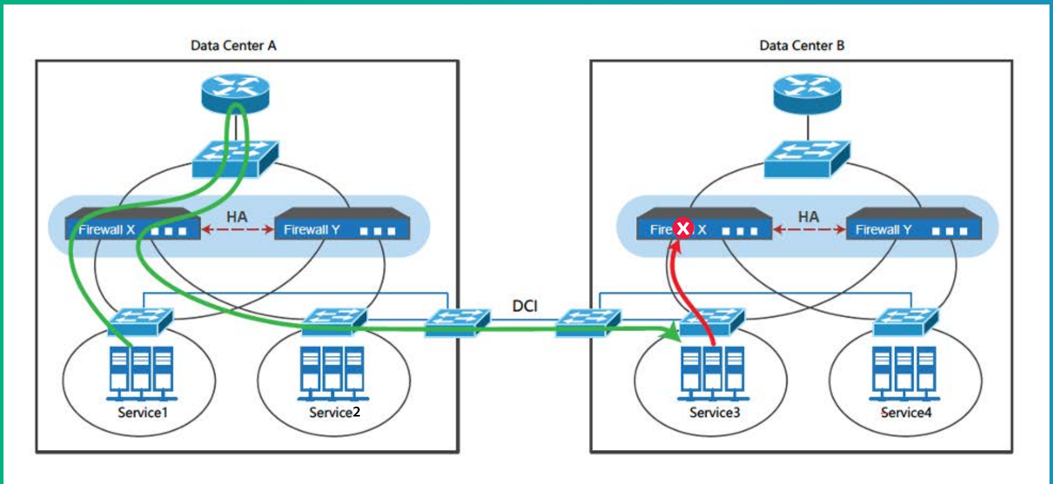

Figure 1: How this issue can occur in a redundant data center scenario.

© 2021 Hillstone Networks All Reght Reserved

Hillstone's Firewall Twin-Mode solves the asymmetric flow issue by synchronizing both pairs of redundant data center firewalls via dedicated data control links, essentially creating a single logical firewall. .

- A transaction sends traffic from servers in Data Center A (on the left) to servers in Data Center B (right).
- IT policy requires all Data Center A traffic to traverse Data Center A firewalls (see green line).
- Instantly, Twin-Mode Link synchronizes the Data Center A Firewall session configuration and state information with Data Center B firewalls (Red dotted line labeled Twin-Mode).
- The return flow hits the Data Center B firewalls, which, thanks to Twin-Mode, are aware of the session established on firewalls in Data Center A. They forward the return flow to Data Center A firewalls.
- The return flow passes through the Data Center A firewalls successfully and completes the transaction (purple line), providing access to the requested information.

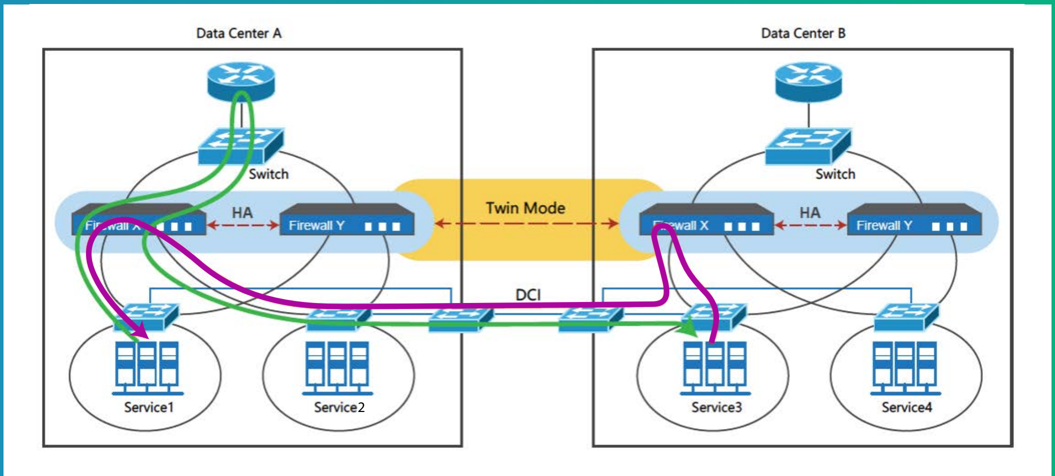

Figure 2: Illustrates how Firewall Twin-Mode would work in the same redundant data center scenario described above.

© 2021 Hillstone Networks All Reght Reserved

# Secure Business Continuity Achieved

Hillstone's Firewall Twin-Mode offers enterprises several important business benefits. Thanks to Twin-Mode, enterprises can:

- Automatically synchronize firewall configuration and session information among two sets of firewalls in redundant data centers. Stateful firewall failover can occur to a second data center to ensure business continuity under the most trying conditions.
- Provide security to asymmetric flows across data centers.
- Achieve not only security, but full visibility into all data-center-to-data-center traffic across DCI links, thanks to Hillstone's security management platform interface. Previously, DCI traffic was often invisible to data center administrators.
- Achieve both full security and 24/7/365 business continuity across a large variety of data center high availability architectures.

High availability data center design is critical for businesses that rely on continuous system availability. Hillstone's Next-Generation Firewalls offer Twin-Mode to enable the redundant data center design that can achieve 24/7/365 availability without compromising security.

© 2021 Hillstone Networks All Reght Reserved

Hillstone NETWORKS

Visit **www.hillstonenet.com** to learn more

or contact Hillstone at **inquiry@hillstonenet.com**