F5 White Paper

Management Networks—Living Outside of Production

Management networks segregate non-production traffic off production networks.

by Paul Stalvig
Technical Marketing Manger

Contents

Introduction

Creating a global personal or business presence today is only a few clicks away, thanks to the new capabilities in Information Technology (IT). Data communications has also changed dramatically. Many of today's applications are housed in large data centers; because of that, they are distributed to remote users who operate these applications away from the central data center. In addition to the usual contingent of business applications, we now have applications that provide numerous real-time or near real-time information and enable real-time transactions. One example is NOAA's weather widgets that give the latest weather updates from the desktop or web browser. Another example is the many music, video, and television programs available for download.

This activity has changed the character of the Internet and how we use it. Today's Internet is an ever-changing source of information and entertainment, limited only by the imagination of developers. The Internet has grown so much that the original users (academia and government) have been forced to move on to their own secluded world (Internet II).

However, one thing has not changed—the need to manage the network. Network configuration, control and troubleshooting are critical capabilities that are most needed to keep the network up and running. When communications fail, applications tend to fail, as remote applications depend on the network to stay fast and available.

Many organizations choose to run data traffic and network management traffic over the same network links. This paper will discuss why it is critically important to segregate network management and control data traffic to ensure that problems can be solved even when the network is melting.

Networks

Over the years, IT departments have struggled with the need to separate both network and application management (or support) traffic from production network traffic. Many terms have been applied to these various networks—for example, administration, management, overlay, or support networks. Still other terms identify networks that serve a specific purpose—firewall, monitoring, security, and storage networks. Each of these networks provides a way to deliver a service to an administrator.

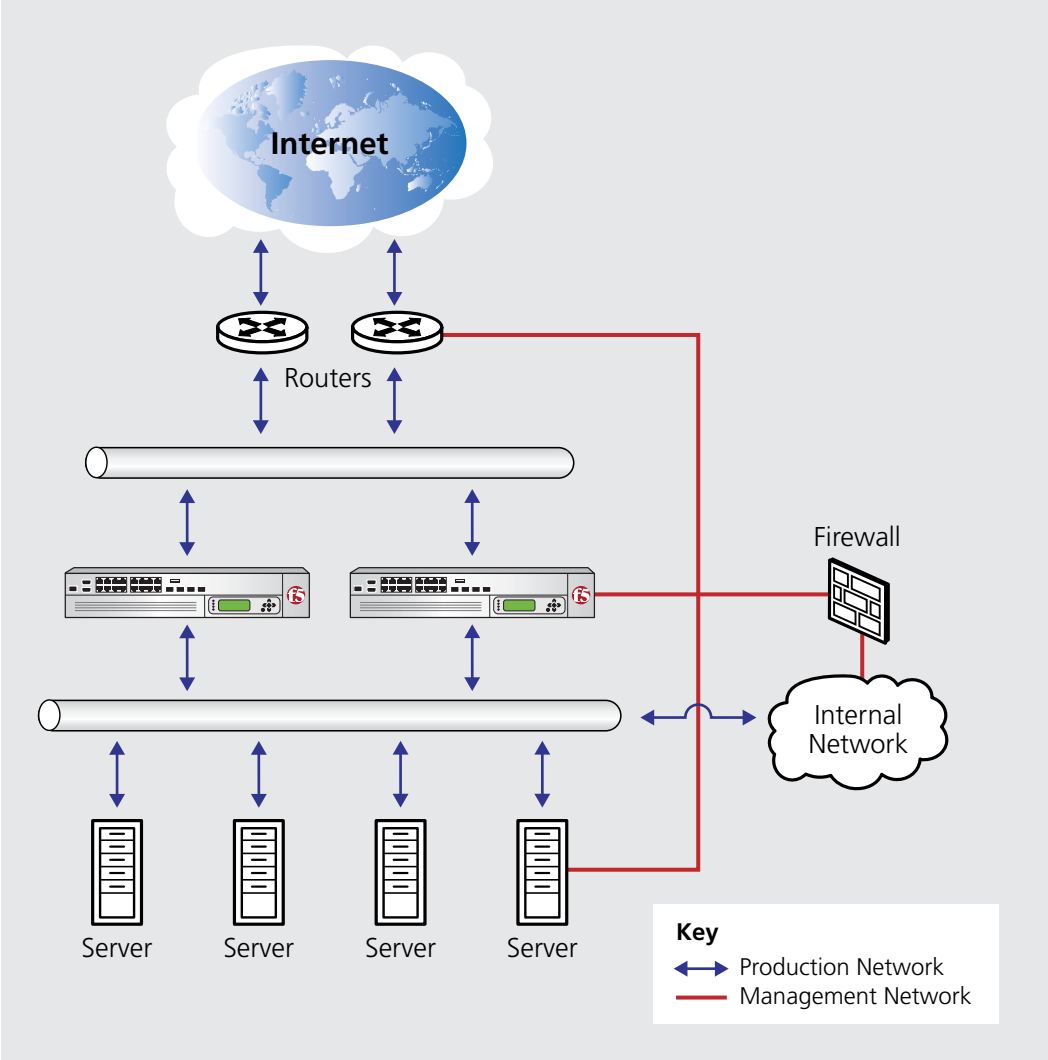

Figure 1: *Management Network*

The problem with using a single, physical path for production and management traffic is that business-related application traffic is forced to compete for the same resources as the traffic that administers, monitors, and supports the network. Management traffic makes up a significant portion of traffic, rising as network devices and servers alert management systems to issues in the network. In fact, management traffic can significantly impact application traffic, using approximately 5–10 percent of available bandwidth during normal operating conditions depending on the amount of monitoring configured for the network. When problems are present within the network, management traffic and production traffic can increase to a level that significantly degrades or interrupts management and application traffic. Connecting to and repairing a problem on an offending device requires the administrator to use this same, degraded network to attempt to fix the problem.

Dynamic routing protocols—RIP, OSPF, and so on—enable routing tables to change as the possible routes change. They are usually standard in many organizations.

While dynamic routing protocols are required in larger networks, they can cause significant problems when errors occur. The downside of their implementation is that a single route can begin to oscillate, or "flap" between two available network paths. This is typically caused when a routing update or router heartbeat signal fails to reach a router. During a route flap, routers inform each other of the loss of a route, and then a new route is calculated and implemented. Production traffic is often interrupted during this time, yet packets continue to enter routers from upstream networks. The routers buffer traffic for as long as possible, but packet loss is almost inevitable as buffers are exhausted.

When route opens back up, the buffered packets are sent through the new link. However, packet loss causes the receiving node to request that packets be re-sent. After a dropped packet, all packets within one window may have to be re-sent to the receiving node; this means that after the packet is dropped, all packets within that window only to clog the network. They will never be used. It may be the case that the route flap continues, exacerbating and perpetuating the problem

The combination of the route flap, the traffic spike that results as router buffers are emptied, and the traffic that must be re-sent due to packet loss can congest the network and prevent management traffic from reaching its destination. By using a management network, administrators are able to immediately access and rectify the issues.

Even with tight budgets and other economic pressures, network managers must use production networks to their fullest. However, even a busy network with no problems can still become congested by increasing production traffic to the point where it can no longer be effectively managed.

Performance and Monitoring

Segregating production and non-production networks eliminates the contention between them, especially if a production link or device becomes congested or experiences a route flap. First, it ensures that the traffic monitored on production links is primarily application traffic. Second, it makes certain that traffic and application monitoring can continue even when links or devices approach their limits.

Knowing the state of your network is extremely useful for capacity planning. If you consistently experience traffic spikes at certain times of the day, you may be able to reschedule tasks such as backups or de-prioritize traffic such as email to ensure that business-critical traffic has sustained and adequate bandwidth. Seeing a trend of increasing network traffic can help you understand when

to raise bandwidth capacity or add other capabilities such as acceleration to your network. The trend would show only the production traffic, since the management and support would use the management network. These actions can help save money and make you look like a hero to upper management.

Monitoring should always be part of your network and application planning. Monitoring helps you understand the "who," "what," "where," and "when" of your production traffic. If the production network is hampered by too much management traffic, monitoring results will be skewed and you may find it far more difficult to determine where an issue lies when it is time to troubleshoot.

Monitoring and alerting traffic should be as free as possible from any production network. If a production link or device between a server and the network management system fails and another device sends an alert through the production network, how can it possibly be received by the network monitoring system? However, the alert is much more likely to be received if it is passed through a separate network.

Although most devices offer a management port that can easily be connected to an out-of-band network, others may not have an extra management port or may be designed for monitoring only out of a production port. In this case, the device can be routed to a separate subnet or VLAN, minimizing the effects on the production network and enabling most management traffic to flow via the management network.

One exception to the above is that passive monitoring needs to follow the same physical path as production. The passive nature of the information becomes a measurement of how the actual production traffic is performing. The second exception is the connectivity done between two devices that must connect via their primary production interface.

Security

By using an administration, management, or support network, administrators are often required to login before accessing the network. This provides two critical pieces of security. First, it keeps passwords from travelling on the same network as production traffic, making it harder for people to troll production networks in search of passwords and other information. Second, it can provide a single point architecturally to provide the detailed access registry of who has accessed the network and for what purpose.

In addition, log information can be routed and stored off the systems before anyone can hack into them and remove traces of the nefarious deeds. In other words with the remote logging capabilities, the logs can be stored locally and remotely. By storing them remotely, it would prevent erasing or modifying the remote logs. Nevertheless, this traffic should not be passed alongside the production traffic.

Conclusion

Addressing this issue is a relatively simple task. By analyzing all network traffic, you can determine what traffic is production. Anything other than production traffic should be evaluated for transport on a separate, management network. Deploying and using management networks correctly, can increase the performance and security of production networks.

F5 Networks, Inc.
Corporate Headquarters
401 Elliott Avenue West
Seattle, WA 98119
+1-206-272-5555 Phone
(888) 88BIGIP Toll-free
+1-206-272-5556 Fax
www.f5.com
info@f5.com

F5 Networks
Asia-Pacific
+65-6533-6103 Phone
+65-6533-6106 Fax
info.asia@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
+44 (0) 1932 582 000 Phone
+44 (0) 1932 582 001 Fax
emeainfo@f5.com

F5 Networks
Japan K.K.
+81-3-5114-3200 Phone
+81-3-5114-3201 Fax
info@f5networks.co.jp

WP-LIVING-OUTSIDE-PRODUCTION 07/08