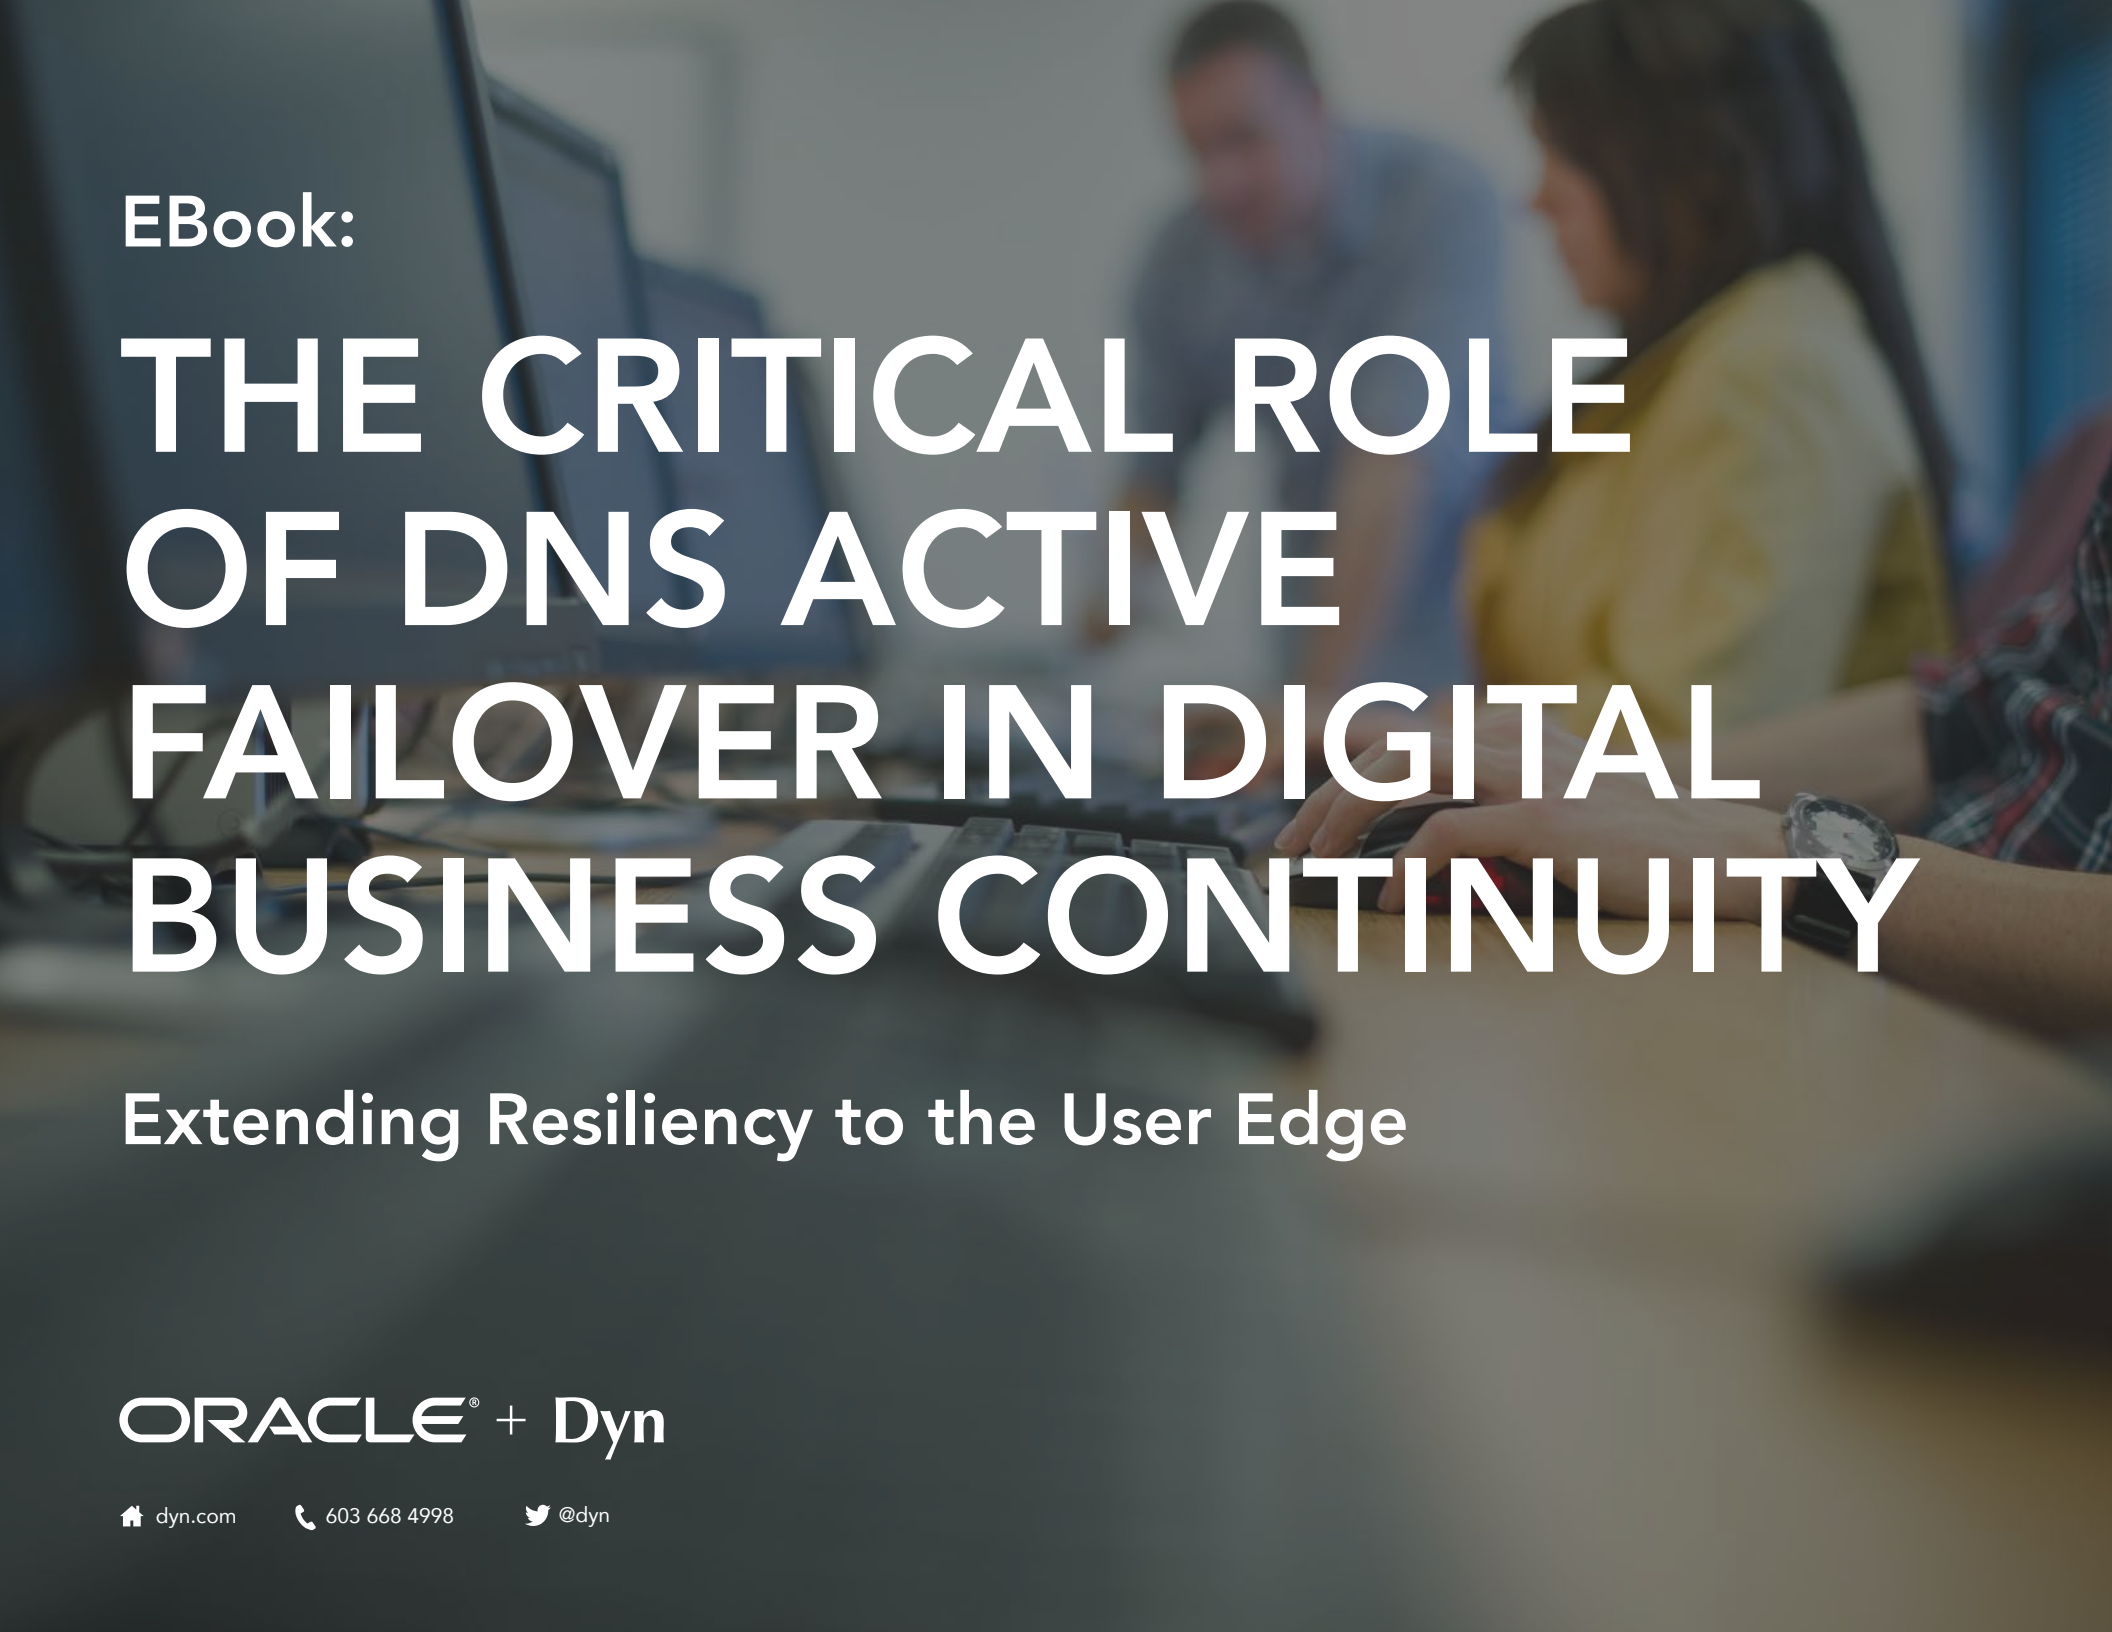

EBook:

THE CRITICAL ROLE OF DNS ACTIVE FAILOVER IN DIGITAL BUSINESS CONTINUITY

Extending Resiliency to the User Edge

ORACLE® + Dyn

dyn.com 603 668 4998 @dyn

The Critical Role of DNS Active Failover in Digital Business Continuity

Extending Resiliency to the User Edge

Downtime in the Digital Age

Digital channels are critical to your business, yet there are a shocking number of digital business disruptions occurring daily. According to a recent Aberdeen study, seventy-eight (78) percent of organizations reported have at least four website disruptions a month, and fifteen percent experience ten or more.[1] Sixty-five (65) percent of organizations in that same study say that it takes them over an hour to resolve an issue.[2]

So, when websites are not available, the cost is high. Reliability—or lack of it—can have a direct impact on the bottom line in a digital economy. According to an ITIC research study:

- 81% of respondents indicated that 60 minutes of downtime costs their business over $300,000
- 33% of those enterprises reported that one hour of downtime costs their firms up to $5 million.[3]

---

1 AberdeenGroup, "How Internet Traffic Steering Is Shaping the Hybrid Cloud Edge," p. 4, June 2017.

2 AberdeenGroup, "How Internet Traffic Steering Is Shaping the Hybrid Cloud Edge," p. 6, June 2017.

3 Information Technology Intelligence Consulting Research, August 2, 2016.

The impact of slow access or downtime on your business is enormous: customer loss, brand reputation damage, or permanent loss of revenue. As more interactions with customer, partner and employees happen over digital channels, an internet failover architecture must be part of your digital business continuity[4] strategy.

Active Failover Is Critical to Business Continuity

Active failover is a DNS service that moves traffic to a healthy endpoint host in the event of degraded service. In such cases, active failover enables your website or web-based applications to remain reachable, provided there's a host available to fail over to.

When the system detects an outage, traffic is automatically rerouted to an alternate, predefined endpoint (this can be done with multiple endpoints in succession). It ensures your traffic finds a route to a healthy location as quickly as possible.

> "According to the Aberdeen Group, a one-second delay in page load time equals 11 percent fewer page views and a 16 percent decrease in customer satisfaction."
>
> *– Aberdeen Group* [5]

4 NOTE TO READER: Digital business continuity, as discussed in this paper, specifically focuses on DNS-based failover strategies at the user edge.

5 AberdeenGroup, "The Performance of Web Applications," p. 4, November 2008, reprinted 2015.

Inside Active Failover

Active failover constantly monitors your primary IP from multiple separate locations. You select the IP address to be monitored, along with the monitoring interval, and the time to live (TTL). For your most vital services, you can set monitoring intervals as low as one minute and TTLs as low as 30 seconds.

Active failover is configured to check on service endpoint health by running HTTP, HTTPS, Ping, SMTP, TCP protocols to verify that the site is still responding. When the service fails to respond from at least two different monitoring locations, your traffic will be redirected to an alternate endpoint. Active Failover considers both the endpoint's ability to serve the user and the condition of the path used to reach that endpoint.

You can also configure the active failover service to send out a notification detailing the status change. It will automatically fail back once the primary is responding again, unless configured for manual failover.

This approach can reduce downtime and accelerate recovery time, with no additional hardware, software, or resources. Properly configured with redundant endpoints in place, no manual intervention is needed—the failover is fully automatic.

Whether it's to address performance and latency issues, or in the case where you lose a site completely, the same DNS technology will reroute traffic to an available endpoint—assuming an end-to-end redundancy strategy is in place from the edge to the endpoint.

Depending upon how you implement your failover configuration, it may be your IT disaster recovery solution, as well as your high availability solution. Your DNS provider should have PoPs in diverse geographical locations around the world, so if there is a natural disaster, the network is able to heal itself to keep data moving.

Common Deployment Patterns

There are three primary active failover deployment patterns. In each of the scenarios described below, active failover is set up to monitor the server at the primary site and fail over based upon the configuration that is chosen.

- **Single-zone Redundancy:** In this scenario, a single zone acts as the failover zone for multiple active zones. This is an affordable option, but could be problematic if multiple failures occur simultaneously or if the failover zone is geographically distant from the site of the failure. Users may experience performance degradation, but they should still be able to reach your site.

- **Nested Failover:** In this case, the failovers are sequential. If a failure occurs anywhere within the configuration, the request will be rerouted to the next available location, until it can be resolved.

- **Nested Failover with Load Balancing:** In this architecture, requests will be routed to a load balancer that will continually be (re)routing based on traffic steering best practices. Should one of the endpoints fail, the rerouting will no longer be based solely on performance, but on the health/availability of the endpoint and path interconnection.

Active Failover Use Cases

There are three primary use case scenarios for active failover. In each of the examples below, active failover is set up to monitor the server at the primary site at one-minute intervals with a TTL value of 30 seconds on the IP address of the primary server. Leveraging this configuration ensures that our customers can automatically serve the failover IP address within 90 seconds.

Use Case # 1. Active:Active Failover: When a company's services are critical to their end users' success, they often need to establish service level agreements (SLAs) that specify extremely limited down time. This is particularly true in industries like financial services. These organizations will often choose to implement a second data center to serve as their backup service location. And it's often designed to be "hot" (online and live) in order to respond to service requests rapidly in the event of a primary site failure. In a solution like this one both sites respond to requests, and when either fails, the other continues to service requests.

Use Case #2. Active:Passive Failover: Having multiple data centers for redundancy is quite common and provides for replication of data at different locations. When the primary data center is not available, you can fail over automatically to a second "cold" site. Online retailers that sell product on their website often choose this type of active failover. If the primary data center is not accessible, their customers can't place orders and will likely buy from a competitor. Companies like this need a seamless, automatic failover solution, with the lowest possible TTL. They will also likely require automatic notification whenever a failover occurs.

Use Case #3. ISP Failover: When it's critical for employees to log in via VPN, organizations cannot afford to lose access because an ISP isn't reachable. Having multiple ISPs configured is the first step in addressing this problem. With this in place, you need a mechanism to fail over from the unreachable ISP to the healthy ISP seamlessly. With an active failover service, customers can maintain uptime and VPN access, no matter which ISP fails.

Conclusion:
Before Bad Things Happen – Be Prepared

Providing a great user experience is always the goal, and the best way to achieve that is by having a well-thought-out digital business continuity strategy. You can't always know what type of disruption you'll face next, but you can be sure that there will be one. It may come in the form of a broken connection but, even more likely, the availability of the application or host. DNS active failover ensures real-time failover to healthy endpoints, allowing you to extend your business continuity solution to the user edge.

Learn more about how DNS-based active failover can enhance your digital business continuity strategy.

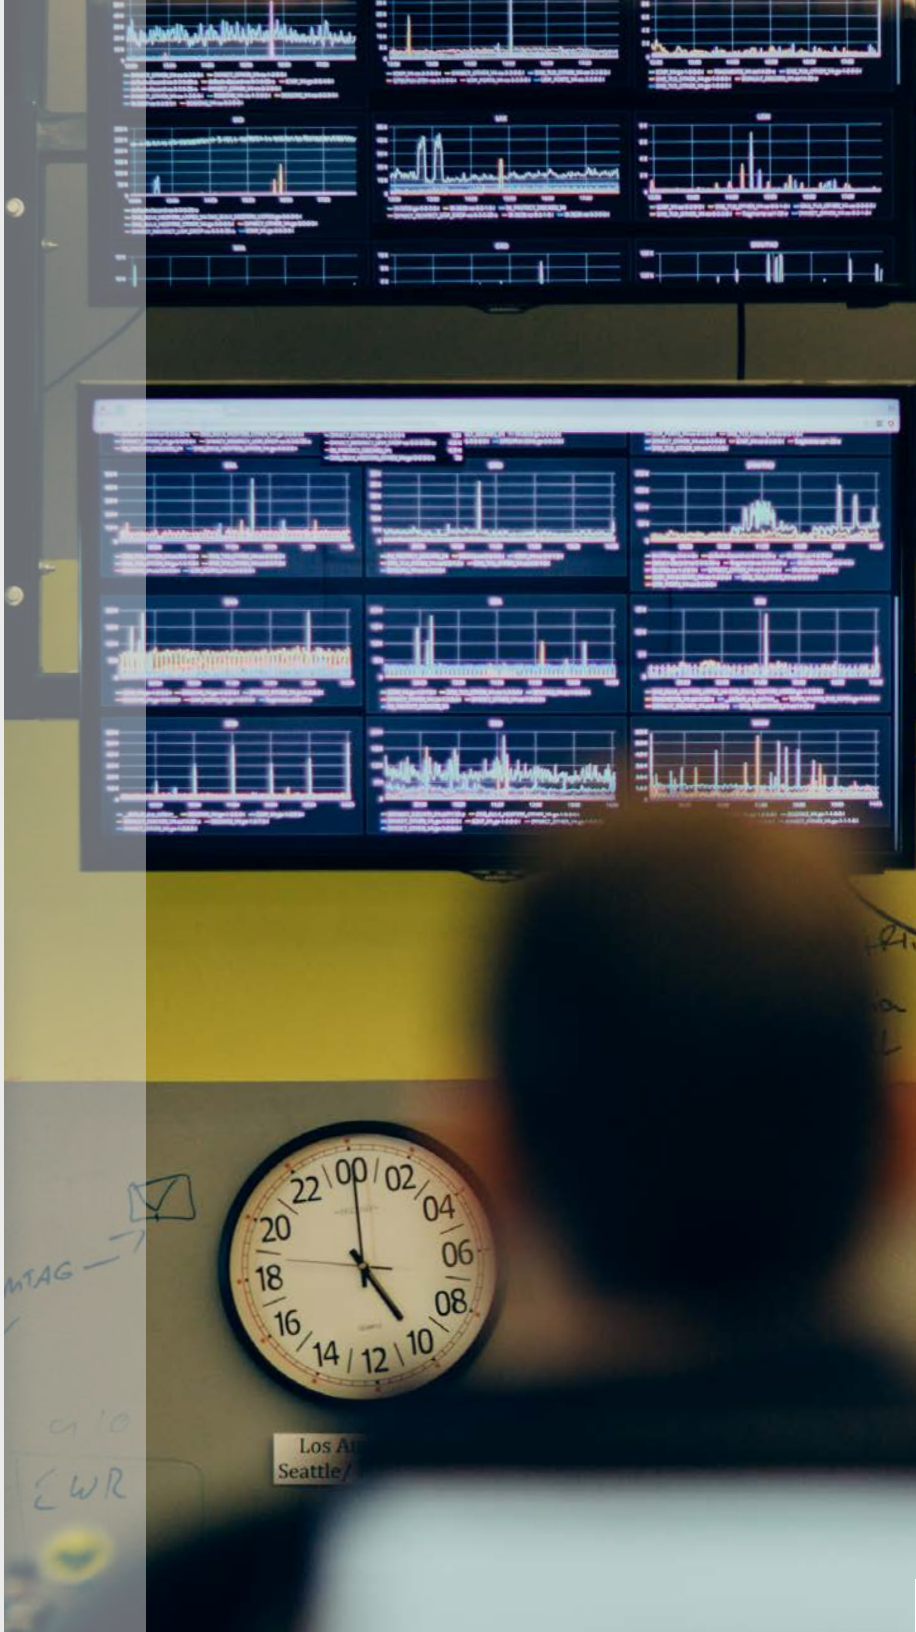

Rethink DNS.

Oracle Dyn is global business unit (GBU) focused on critical cloud infrastructure. Dyn is a pioneer in DNS and a leader in cloud-based infrastructure that connects users with digital content and experiences across a global internet. Dyn's solution is powered by a global network that drives 40 billion traffic optimization decisions daily for more than 3,500 enterprise customers, including preeminent digital brands such as Netflix, Twitter, LinkedIn and CNBC. Adding Dyn's best-in-class DNS and email services extend the Oracle cloud computing platform and provides enterprise customers with a one-stop shop for infrastructure as a service (IaaS) and platform as a service (PaaS).

ORACLE® + Dyn

dyn.com 603 668 4998 @dyn