Proceedings of the 2011 Winter Simulation Conference
S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. Fu, eds.

SURVIVABILITY OF DUAL-CORE NETWORKS DURING RARE EVENTS

Steven Gordon
David Garbin

Noblis
3150 Fairview Park Drive
Falls Church, VA 22042, USA

ABSTRACT

Telecommunication networks are evolving to become more reliable, but many networks remain vulnerable to widespread systemic failures. Reliability of individual components has improved and some networks can achieve availabilities on the order of 0.99999. However, the routing technologies used by these networks, like Open Shortest Path First and Border Gateway Protocol can create system-wide vulnerabilities. Some of the vulnerabilities include widespread outages such as earthquakes and floods, unintentional device mis-configurations, and hacker attacks. One of the leading-edge architectures to address system-wide outages is the use of a dual-core backbone, which uses two independent long-haul cores to connect the network's sites. The network is designed to tolerate the failure of a single core and leave the network fully functional. This work presents an OPNET simulation model of a dual-core architecture. This model predicts the restoral time of various network outages for different device configuration options and different topology options.

1 INTRODUCTION

Reliability and availability are becoming increasingly important in enterprise communications networks. More and more, the loss of enterprise network services translates into loss of revenue. The outage on April 21, 2011 of Amazon's Web Services cloud is an example. This outage denied services to many online sites, including Reddit, HootSuite, Foursquare and Quora. (Pepitone, 2011)

Partial outages caused by infrastructure failures are generally less frequent and less service-affecting than past generations of networks. Equipment reliability and management practices have evolved to minimize the effects of random outages. As a result, network availabilities of 0.99999 have become achievable goals.

However, the network and management practices that produce these availabilies often do not consider the events that can create widespread, systemic outages as opposed to limited service outages, which is typically the case in random failures. These system outages often have root causes that include floods, hacker attacks, terrorist attacks and unintentional device mis-configurations. In addition, the results of these systemic outages are increasingly severe, as in the case of Amazon in their loss of delivery of web services to their customers.

In some of these circumstances, the backbone transport component of the enterprise network is the target of the issue. When viewed as a single element, the backbone transport can become a single point of failure. In order to address this issue, a leading-edge architecture technique is the creation of a *dual-core* backbone, in which two separate and independent cores simultaneously provide transport for all of the network's sites. This dual-core architecture allows the network to remain functional and fully-capable in the event of one core becoming inoperable. In addition, this dual-core allows the network managers to shut down a malfunctioning core while leaving the network in a fully operational state.

This paper presents the basics of a dual-core architecture and a simulation model that was created to predict the performance of a dual-core network when subjected to various events. The first section overviews the basic elements of a dual-core architecture. The second presents the simulation models. The third details the experiments performed and the outputs produced. Finally, the last section provides a summary and wrap-up.

2 BACKGROUND

2.1 Basic Dual-Core Architecture

In a dual-core architecture, the subnetworks that terminate user access lines (*distribution networks*) are viewed as distinct subsystems and are connected through IP routing. The long-haul transport component (*core*), is also conceptualized as a separate subsystem (Figure 1). Each subsystem (cores and distribution networks) has its own internal routing architecture that routes within the subsystem, and uses Border Gateway Protocol (BGP) to route packets to neighboring subsystems to form end-to-end (ETE) connectivity. This architecture creates a separation of subsystems with BGP boundaries so that (1) faults within a subsystem can be contained within the BGP boundary, leaving the outside subsystems insulated from the faults, and (2) each distribution network can use the extensive set of path-selection mechanisms in BGP to choose the best core for outbound traffic.

The objective of the dual-core architecture is to create a paradigm where the core subsystem can be

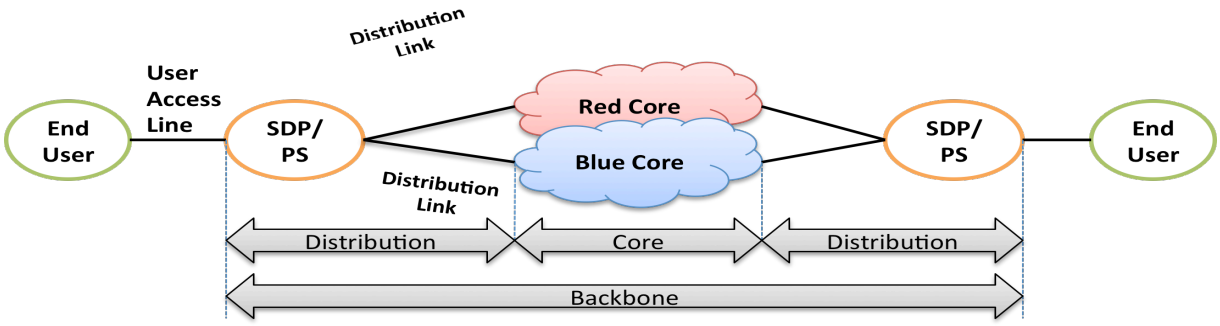

Figure 1: Notional Architecture

replicated and two core subsystems can be used in a parallel, redundant manner. To realize this, the core subsystems must be completely independent, and the end-user subsystems must be able to use either or both of the core subsystems. *Independence* means that the two cores must be separate at Open Systems Interconnection (OSI) layers 1, 2 and 3. Specifically, they share no physical circuits, switches or routers. In addition, the layer 3 routing processes must be *distinct*, and the cores must not communicate with each other.

However, in order to allow customer's traffic to use either core, a path selector (PS) device must be able to route to either core, and the PS therefore will be the only device that can "hear" routes from both cores. Because the cores physically join at this device, the routing architecture prevents routes from the cores to be mixed in the PS. In addition, the PS must have the functionality to make the core-selection decision on a desired criterion, which can include, but is not limited to the following:

- Manual request for a core shutdown
- Manual request for routing to use a specified core
- Core that is unable to reach the needed destination(s)
- Core that is experiencing high latency or packet loss.

2.2 Study Objectives

The sponsors of this work undertook this work to study the operation and possible options with the dual-core architecture. Because dual-cores are not in widespread, common use, their characteristics and options are not well-understood broadly across the communications field. Specifically, the objectives were to

- Verify and predict rerouting times for various types of failures
- Create a tool with which optimization studies could be performed
- Observe the behavior of dual-cores under various scenarios
- Verify that specific configuration options, in particular BGP timer values, will result in stable network operation.

The last item is particularly important. The studies participants expected that stability in BGP would be an issue, and that it would be highly dependent on the specific configurations. Instabilities are difficult to predict in an analytical manner, and this type of discrete-event simulation was deemed to be the best prediction tool of possible BGP issues.

3 SIMULATION MODEL

This model was built in OPNET Modeler, which is a discrete event communications simulator. It simulates the circuits and devices, such as switches and routers on a packet-by-packet basis. The model consists of the typical network elements in the actual network, e.g., routers, switches, and circuits. Each device can have an associated configuration that customizes the devices' operation in a similar manner to an actual network. Each of the devices contains a node model that defines the protocols that govern its actions on the packets. Typical protocols include Ethernet/802.3 (layer 2), IP (layer 3), OSPF (layer 3), and BGP (layer 3).

OPNET simulates the network by simulating the action that each protocol in each device applies to each packet. All packets are simulated individually, and the operation of the network as a whole is derived from cumulative effect of the journey of each individual packet. From this, all traffic flows, link utilizations and routing behaviors can be deduced.

Our model seeks to emulate a typical IP service-provider network. Seven sites, which are at least doubly-connected, were chosen to make a small, but representative network. The middle site (CHI) is the only deviation from a typical design. Its connections were chosen so that its failure would cause the network to become disconnected. This was chosen so that failover scenarios can be easily constructed (Figure 2).

3.1 Core Architecture

Each core is configured with OSPF as the interior gateway protocol (IGP). This routing protocol gives the core its recovery mechanism from link and node failures that do not cause the core to become disconnected. Convergence of OSPF is typically fast, but routing for any or all of the routers can change upon OSPF convergence.

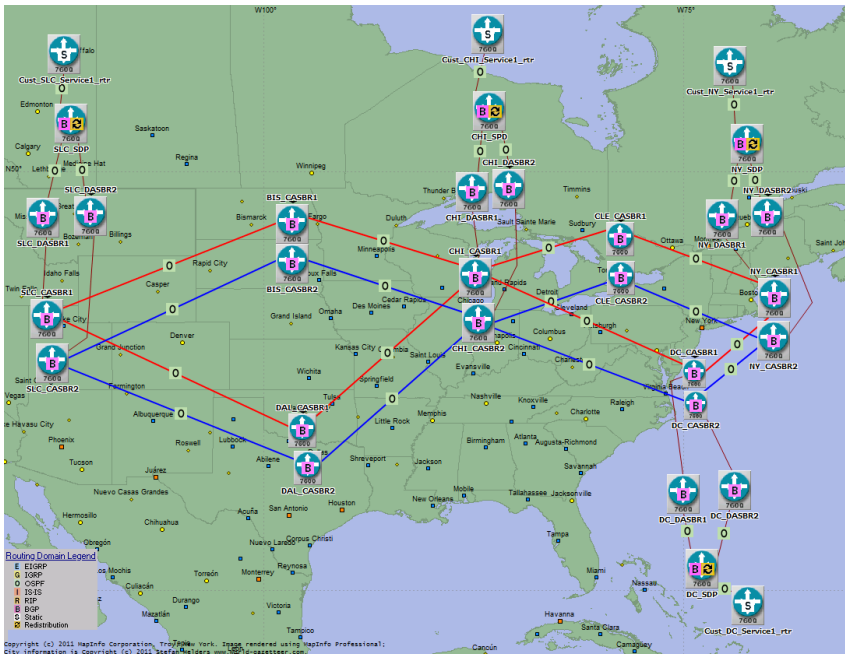

Figure 2: Physical Topology

Each core is also configured with BGP, which is the exterior gateway protocol (EGP). The EGP defines and restricts the routing of the traffic between subsystems (distribution and core networks). In general, BGP implements "policy routing," which allows networks from different political entities (enterprises) to exchange traffic, but gives each entity the ability to decide the types of traffic to exchange. A single BGP domain, *Autonomous System (AS)*, is subject to a set of policies created by a single political entity. However, one entity may operate multiple ASs.

The edge of the core, which is the boundary to the customer ASs, is demarked by a core AS boundary router (ASBR). The core ASBR:

- Runs External BGP (EBGP) to exchange routes with the distribution ASBR (DASBR), which is the boundary router in the distribution AS. This exchange allows (1) originating traffic from the customers to enter the core and get routed towards its destination, and (2) terminating traffic in the core to get routed to the proper distribution network towards its destination.
- Runs Internal BGP (IBGP) to exchange BGP routes across the core, so that each ASBR can learn a full set of customer routes from all of the customer ASs.

The core ASBRs (CASBRs) are configured with a full-mesh of peer relationships. A peering configuration allows the peering nodes to exchange BGP routes, so that each BGP peer will be able to reach all of the destinations learned by all of its peers (Figure 3). Each of the two cores will have a full mesh of IBGP peering **within its own core.** In this manner, each core ASBR of each core will be able to reach all customer sites without using the other core.

3.2 Distribution Network Architecture

The distribution network is the segment of the network starting from the distribution ASBRs and extending back to the service delivery point/path selector (SDP/PS) device. Each distribution network, like the core, has its own IGP (often OSPF) and instance of BGP. Note that each of the distribution networks runs a different instance of the IGP and BGP. As a result, any instability in the IGP or BGP routing process in any particular subsystem is not transmitted to the neighboring subsystems.

The distribution ASBRs peer with core ASBRs. In the model, the nodes designated with "DASBR1" peer with the nodes designated with "CASBR1", and similarly for "DASBR2" and "CASBR2". The distributions ASBRs do the following:

- Announce the local routes to the core, so that inbound traffic can be properly routed to the destination AS and host,
- Listen to the EBGP announcements from the core ASBR to learn the path for the outbound traffic from the distant distribution networks.

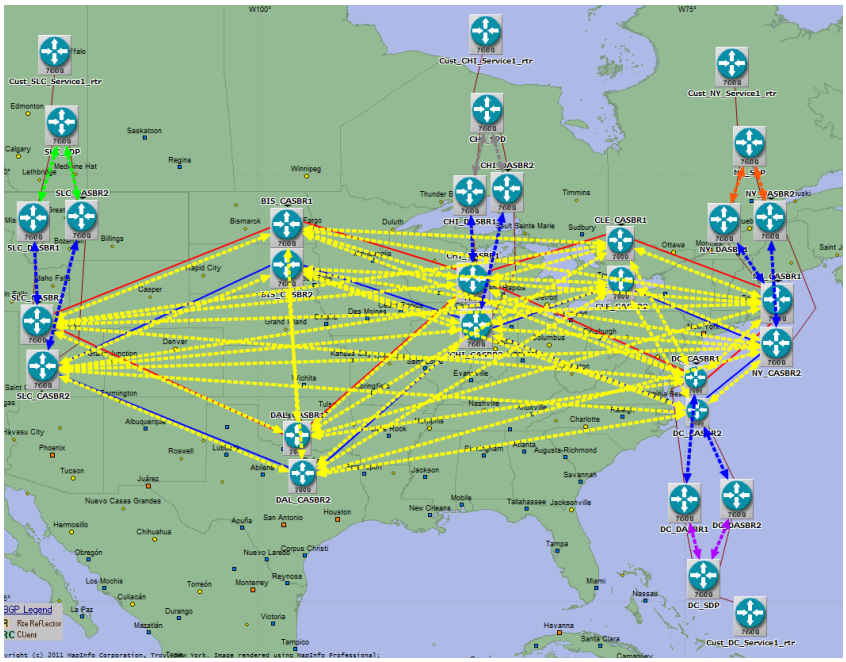

Figure 3: IBGP Peering (Dashed Yellow)

In a fully operational network, each of the two distribution ASBRs in each site hears an identical set of BGP routes from each of the core ASBRs. That makes each distribution ASBR a viable transit point for backbone traffic. Figure 3 shows the EBGP peers in dashed blue.

The PS device must perform the core-selection function, and requires the routes from each core to support this decision. To support this, PS device is configured to run BGP, and an IBGP peering session is configured from each PS to both distribution ASBRs (Figure 3, green, red and purple lines). With this information, the PS can make the following types of selections:

- **Manual selection**, by configuring a preference for either blue or red routes. This can be done on a per-source, per-destination or per-application basis.
- **BGP default**, where a list of BGP route selection criteria are evaluated to find the "best" route. In this configuration, the core selection tends to be arbitrary as long as both cores are operational. If one of the cores is non-operational, its routes will be removed from the PS device's routing table, and the remaining core will become the only choice.
- **Performance-based**, where the PS would employ an active network monitoring tool to determining the reachability state, and/or delay, and/or packet loss to each destination. The results of these active monitoring tests can influence the preferred routes.

3.3 Failure/Failover Requirements

Each core is doubly-connected so that a single link failure will be recoverable on the same core. Hence, for single link failures, the traffic on the impaired core should stay on that core and the rerouting should be performed by the IGP. This should yield fast recovery times for single link failures, which are the vast majority of the failure events.

However, a double link failure or node failure can cause a core to become disconnected. In this state, some or all of the network's destinations may be unreachable from some of the origins. This type of failure should trigger a core failover, in which some of the origins decide to use the other operable core instead of the impaired one. This routing decision can be made separately for each origin-destination pair. For example, if the failure of CHI_CASBR1 leaves the red core disconnected so that SLC cannot reach DC on the red core, but can reach DAL on the red core, then SLC should reroute DC-bound traffic over to the blue core. For destination sites on the red core that are still reachable, rerouting to the blue core is optional. Depending on the Operations policy of the enterprise, either leaving the reachable destination on the red core or performing a rerouting may be preferable. This type of core-failover rerouting would be performed by BGP, and the reconvergence time tends to be much longer than that of IGP rerouting.

4 EXPERIMENTATION

The following section presents the three modeling scenarios:

1. Successive link failures in a core
2. Node failure and BGP parameter tuning
3. Combined-core routing.

4.1 Experiment 1: Successive Link Failures in a Core

In this scenario, first a link fails in the red core that causes the IGP to reroute some traffic and reconverge. Then, 100 seconds later after the network has reconverged, another link failure causes a site (SLC) to become disconnected and should trigger a failover to the blue core. The restoral time is observed for each event for an ETE flow. The restoral time generally consists of two parts:

- Time for origin to make the rerouting decision
- Time for routing changes to get propagated to the rest of the network.

Both of these steps are required to re-establish an ETE flow. Before the failure, the path from SLC to DC is shown in Figure 4.

The first failure was the BIS_CASBR1-CHI_CASBR1 on the red core at 400 sec. After 5.7 seconds, the flow was re-established as shown in Figure 5. At 500 seconds, the SLC_CASBR1-DAL_CASBR1 link failed causing SLC to be isolated on the red core, but remaining connected on the blue core. After about 150.7 seconds, traffic was re-established on the blue core as shown in Figure 6.

The conclusions from this are the following:

- Rerouting of single link failures works as intended
- Rerouting traffic to the other core from origins that become disconnected works as intended
- Rerouting to the other core takes ***much*** longer than rerouting on a single core (BGP vs. IGP rerouting).

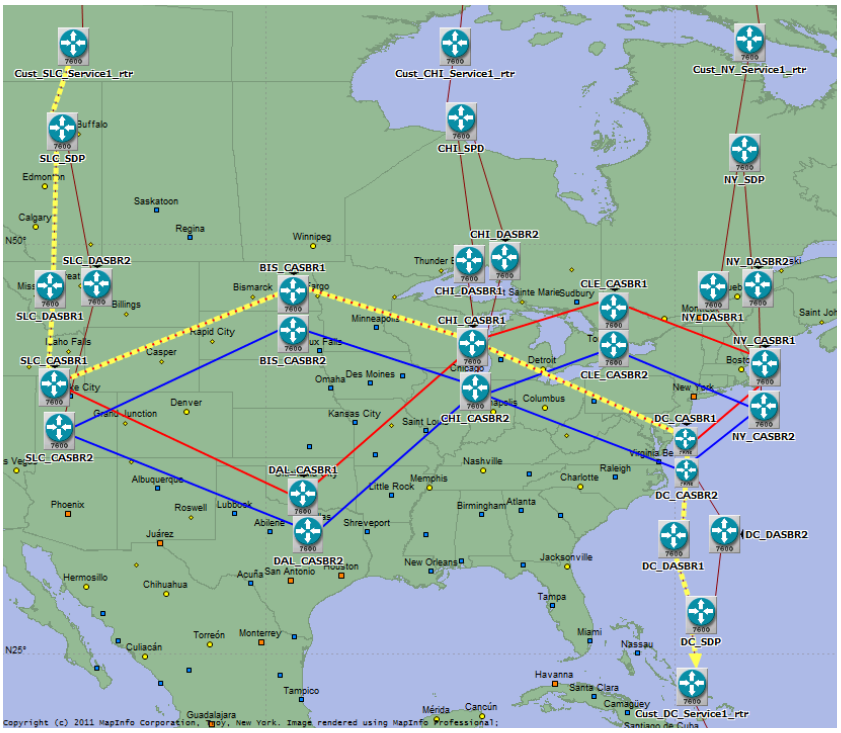

Figure 4: SLC-DC Path

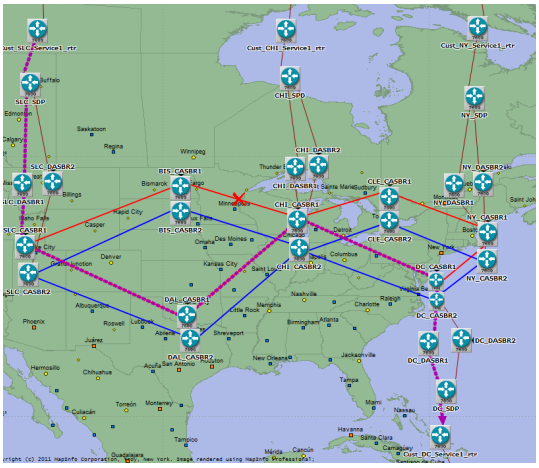

Figure 5: Rerouting on Red Core

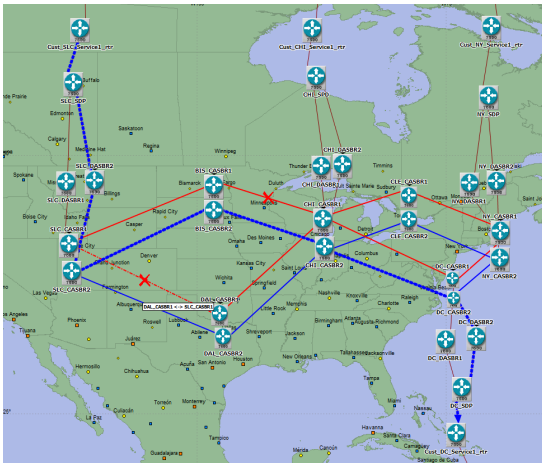

Figure 6: Failover to Blue Core

4.2 Experiment 2: Core Segregation and BGP Parameter Tuning

This scenario consists of two steps. This first is an investigation of the sudden segregation of a core with a single event. As in Experiment 1, the restoral time is observed. Node CHI_CASBR1 (red core) was failed. The configuration of the model specified that the device failed, but the links remained active at layer 1, which will leave the remote ends of the links unaware that the node had failed. This configuration was done to ensure that no link outage detection mechanism is initiated to speed restoral time, and

outage detection can only be done by BGP though the loss of BGP KEEP-ALIVE messages. The resulting restoral time will be compared to the second step.

In step two, the scenario is rerun with improved BGP parameters. In particular, the following parameters are changed:

- Hold timer: 180 sec (Cisco default) changed to 60 sec.
- Keep alive timer: 60 sec (Cisco default) changed to 20 sec.

Again, the restoral time is observed and will be compared to step one.

Step one begins with the flow shown in Figure 4. After failure of CHI_CASBR1, the flow is rerouted to the blue network as shown in Figure 7 with a restoral time of 138 seconds.

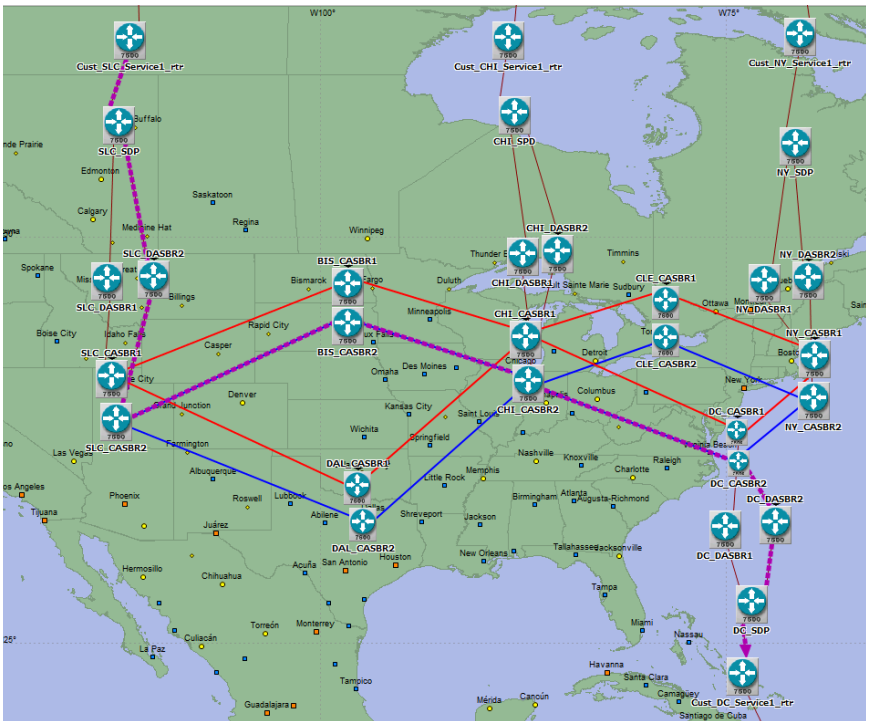

Figure 7: Rerouting After CHI_CASBR1 Node Failure

In step two, the same experiment is performed with the improved BGP timer values. The same rerouting to the blue core results, however, the failover time is much improved. Figure 8 shows the traffic (bits/sec) from the SLC-DC flow for the default BGP timers in black and the improved timers in pink. Both traces drop to zero at 500 seconds, when the CHI_CASBR1 router fails. However, the simulation yielded a restoral time of 138 seconds with the default parameters, vs. 43 seconds with the improved parameters.

Although the results of this experiment yielded an expected result, i.e., a decreased reconvergence time, it also shows that the dual-core remains in a stable mode of operation with this timer improvement. One of the issues with decreasing these timers and changing BGP behavior is that, at low enough settings, a "thrashing" effect is likely to result where BGP neighbors are incorrectly assumed to have failed by their neighbors, and a network reconvergence is needlessly initiated by BGP. This simulation also found that this behavior was avoided at these timer values.

4.1 Combined-Core Routing

Combined-core routing is a tactical option that allows network managers to restore ETE service for some failures when one core is disconnected and the other core has inoperative distribution link(s). For some failures, combining both cores into a single IGP and BGP routing domain will restore connectivity. Figure 9 shows this topology, which is the same as the topology shown in Figure 2 with the addition of links between the CASBR1 and CASBR2 routers at each site. Normally, these links are disabled by setting their configuration status in the CASBR1 and CASBR2 routers to "administratively down." When set to "administratively down", this configuration leaves the same IP topology as Figure 2, with two independent IGP and BGP routing domains.

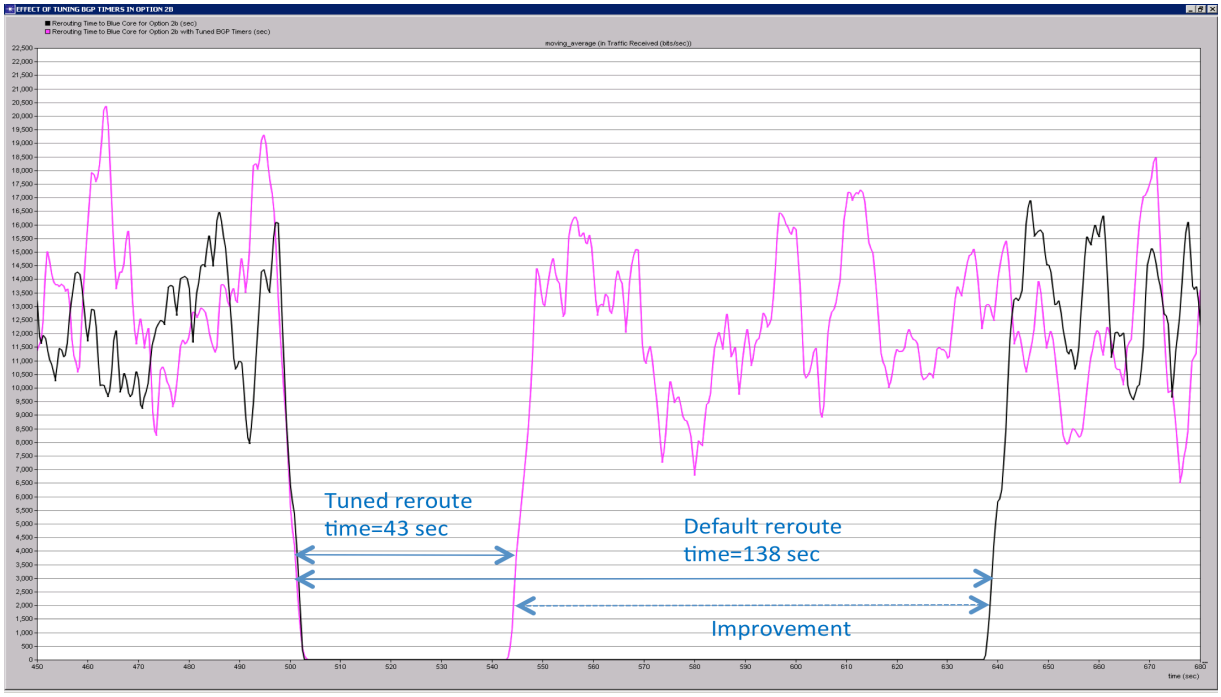

Figure 8: BGP Restoral with Default and Improved Timer Values

However, these cores can be combined into a single IGP and BGP routing domain by changing the configuration of these links to "administratively up" on the router interfaces. This will allow the two IGP and BGP instances to form a single instance. Network operators may want to do this if all of the following conditions apply:

- One core is segregated
- The remaining core has at least one distribution link down
- In the sites that the above distribution line is down, the other one is up.

This scenario tests this capability, and is performed in three steps:

1. At 400 seconds, CHI_CASBR1 fails, which segregates the red core.
2. At 500 seconds, SLC_CASBR2-SLC_DASBR2 fails, which leaves SLC_SDP is isolated from rest of network. However, one access link and one core are still usable.
3. At 800 seconds, all cross links are changed to administratively up and the combined core is formed.

A flow from SLC to DC was studied and its rerouting time was observed. After the network converged after step 3, the routing converged to the path in Figure 10. Figure 10 shows the flow from SLC beginning on the red core, because of its functional distribution line, then hopping to the blue core via the BIS cross-link. After the cross-links were activated at 800 seconds, both core were routing as a single core, and traffic will use cross-links as part of an ETE path. Figure 11 shows the rerouting times for the steps above. After the failure in step 1, Figure 11 shows that traffic was restored in 53 seconds. The failure in step 2 is shown by the trace dropping to zero at 500 seconds. Finally, Figure 11 shows that after step 3, ETE traffic was restored in 64 seconds. Note that the BGP parameters were set to the optimized values in the section, "Experiment 2: Core Segregation and BGP Parameter Tuning."
From this experiment, the following can be concluded:

- Combined-core routing is potentially an option for certain severe network outages.
- The restoral time is on the order of 64 seconds.

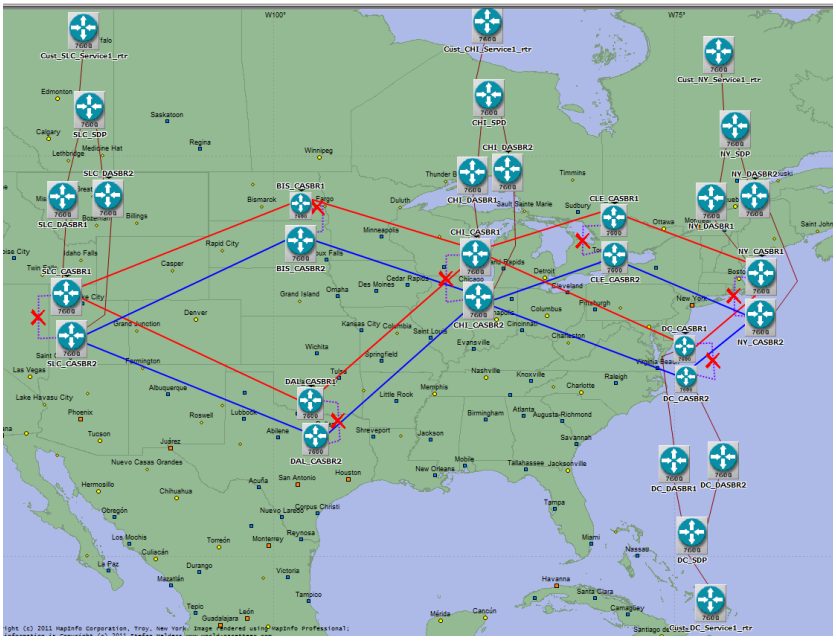

Figure 9: Combined-Core Topology

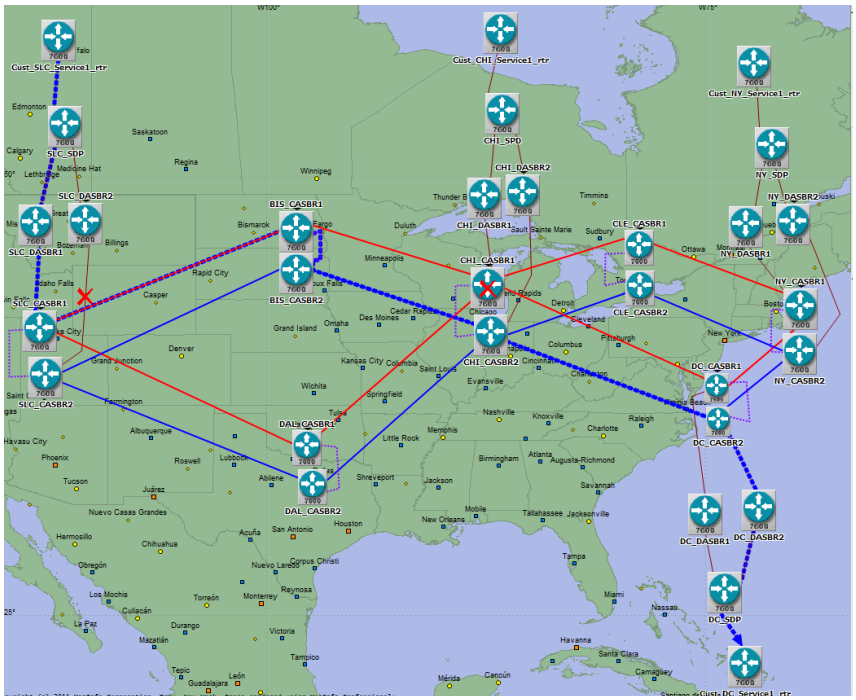

Figure 10: Routing with Combined-Core after Failures

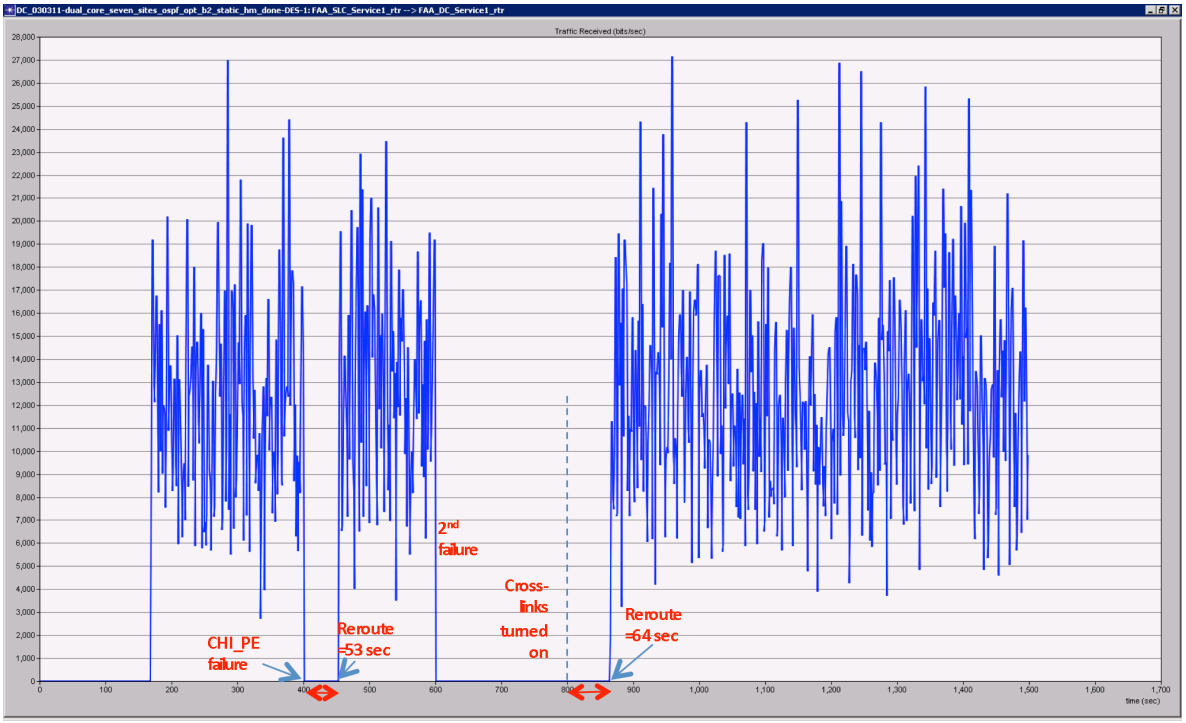

Figure 11: SLC-DC Flow in Combined-Core Network

5 SUMMARY

This type of dual-core architecture in not yet in widespread use, but it has been adopted by some organizations with very high availability or fault tolerance requirements. Some members of the financial community have adopted this architecture to more thoroughly safeguard against outages. For this community of interest, loss of network services for even a short time translates directly into large amounts of lost revenue.

This study was useful to its sponsors to investigate the design and operation of a dual-core enterprise networks. Before performing the study, some of the major research questions included:

- Was the concept of operations viable?
- If so, what performance characteristics could be expected?
- Was the behavior of dual-core networks stable and predictable?

This study resulted in confirming that the dual-core architecture was indeed a viable option, network behavior was stable during failures, and the performance is highly dependent on various detailed configuration options. In pursuit of the goal of improving the performance, the simulation model is an extremely useful tool. It allows different detailed device configurations to be loaded and simulated under a variety of failure conditions, and for the performance to be estimated under those conditions.

REFERENCES

Pepitone, J. 2011. *Amazon EC2 outage downs Reddit, Quora*. Retrieved May 17, 2011, from CNN Money: http://money.cnn.com/2011/04/21/technology/amazon_server_outage/index.htm

AUTHOR BIOGRAPHIES

Steven Gordon is a Principal in the Noblis. He has 24 years in the area of Telecommunications, including 9 years with Telecommunications Service Providers. His areas of expertise include network architecture, network design, capacity planning, capacity management and performance analysis of large networks. His recent work includes performance analysis of applications for the Department of Defense and failure analysis of network infrastructure for the Federal Aviation Administration Mr. Gordon holds a Master of Science in Systems Engineering from Case Western Reserve University.

Mr. David Garbin is a Senior Fellow in the Noblis' Enterprise Services Mission Area. He has 40 years experience in the telecommunications and networking field, focusing on the design and economic analysis of large networks, both for carriers and their customers. He recently served as Vice President of Strategic Network Planning for Cable & Wireless, one of the largest global providers of Internet services. His current duties include research into providing quality-of-service in convergent IP networks and in the use of advanced modeling and simulation to analyze and design resilient networks supporting the nation's critical infrastructure. Mr. Garbin holds advanced degrees in electrical engineering from MIT and is the co-author of the New McGraw-Hill Telecom Factbook, currently published in its second edition.