

How Cisco IT Designed a Data Center that Can be Quickly Transformed from Development to Disaster Recovery

Using data center infrastructure for non-production applications instead of keeping it idle lowers costs and reduces risk.

Cisco IT Case Study / Data Center / Dual-Purpose Data Center: Cisco IT previously backed up the Texas Metro Virtual Data Centers (MVDC) to a disaster recovery (DR) facility in Research Triangle Park, North Carolina, which was approaching space and power limits. Rather than design a traditional DR facility that would remain idle most of the time, Cisco IT designed a data center that could support non-production applications most of the time, and be quickly repurposed for DR. This case study explains how Cisco IT designed a data center to support two environments: non-production and DR. Cisco customers can draw on Cisco IT's real-world experience to decide whether this approach would work for them, and to design their own non-idle disaster recovery facility.

“Organizations struggle to afford a separate infrastructure for DR. Cisco IT realized we could take advantage of virtualization and automation to repurpose the same infrastructure for DR and non-production applications.”

James Cribari, IT Manager, Cisco Global Data Center Program

Challenge

As part of its mission to deliver IT as a Service, Cisco IT Global Data Center Program is migrating critical applications, network, compute, and storage to a private cloud built on the Cisco Unified Computing System™ and Cisco Nexus® Series Switches.

Disaster recovery is a key component of the private-cloud strategy, and Cisco IT needed a

new disaster recovery facility for the company's Texas Metro Virtual Data Center (MVDC) in Richardson and Allen, Texas. The MVDC operates approximately 1300 production applications, including Oracle eBusiness Suite.

Previously, these applications were backed up to Cisco Data Center – Raleigh 5 in Research Triangle Park, North Carolina, which was approaching its limits for space, power, and cooling. “Cisco's overall data center strategy called for a new data center that would increase resiliency, meet future capacity requirements, and support diverse Cisco lines of business,” says James Cribari, IT manager for Cisco's Global Data Center Program.

Rather than designing a traditional data center that would remain idle except during disasters, Cisco IT decided to design a dual-purpose data center. Ordinarily it would be used for non-production applications: development, test, and staging. In the event of a regional disaster affecting the Texas MVDC pair, Cisco IT could initiate automated processes to free up resources for disaster recovery. “Organizations struggle to afford a separate infrastructure for DR,” says Cribari. “Cisco IT realized we could take advantage of virtualization and automation to use the same infrastructure for DR and non-production applications.”

Dual-purpose infrastructure would cost significantly less than dedicating equipment for DR and having it operate in standby mode. “Maintaining an idle architecture for DR wastes space and resources,” says Jag Kahlon, IT architect for Cisco. Putting the infrastructure in service when not needed for DR would also eliminate need for constant testing, a standard practice for dedicated DR infrastructures, which tend to atrophy when idle.

Design goals for the new facility, called Cisco Data Center – Raleigh 1, included:

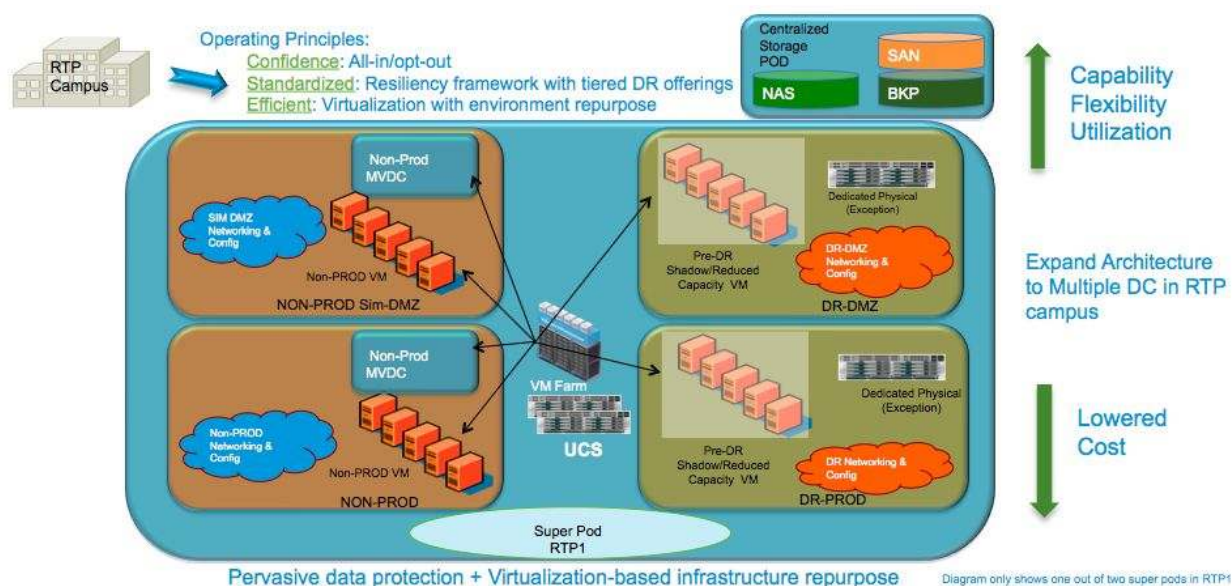
- Repurpose infrastructure from non-production to disaster recovery, using automated processes.
- Virtualize to the fullest extent possible, approximately 80 percent of applications. “To make best use of space, power, and cooling, we wanted to virtualize all non-production applications in Raleigh 1,” says Ramona Jackson, IT manager at Cisco. “This would require server and network design to support applications with high transaction-arrival rates, such as Cisco’s Quote to Cash applications, which touch 100 percent of Cisco sales transactions.”
- Operate sustainably by meeting requirements for U.S. LEED-NC 2.2 Gold certification for construction.

Cisco IT decided to consolidate all regional non-production in Cisco Data Center – Raleigh 1 rather than distribute it in different regional data centers. This approach would avoid the costs of redundant infrastructure, reduce environment refresh complexity, and make sure that the non-production environment mimicked production. “For testing to mimic production, we need to maintain transactional dependencies between applications in the non-production environment,” says Kahlon. “We will simulate latency within the non-production environment to match the natural latency between the two Texas data centers that make up the MVDC pair. The alternative, keeping regionally dispersed non-production environments up-to-date, would require pushing a lot of data over the WAN.”

Solution

Like other Cisco data centers, Raleigh 1 is based on the Cisco Unified Computing System and Cisco Nexus switching architecture. The difference is the superpod design. The building blocks of other Cisco data centers are pods, containing computing and storage for one environment, such as production, non-production, or DR. “Raleigh 1 is the first Cisco data center designed with superpods that can support either of two environments: non-production applications or disaster recovery,” says Kahlon. Figure 1 illustrates the superpod design.

Figure 1. Architecture to Repurpose Development Infrastructure for Disaster Recovery



The enabling technologies for the dual-purpose data center, described in the following paragraphs, are:

- Cisco UCS Manager service profiles, enabling Cisco IT to reconfigure blade servers for disaster recovery with a few clicks.
- The Virtual Routing and Forwarding (VRF) feature of Cisco Nexus 7000 Switches, allowing multiple instances of a routing table to coexist on the same physical router. A system administrator simply enters NX-OS commands to reconnect the Cisco Nexus 7000 Switch pair from the non-production VLANs on the Cisco Unified Computing System to the DR VLANs.
- Cisco Nexus 1000V Switch, enabling reconfiguration of the network / compute interface.

Cisco Unified Computing System Configuration: Different Service Profiles for Development and DR

The Cisco Unified Computing System contains Cisco UCS B200 M2 Blade Servers with 96 GB RAM for the VMware vSphere virtual environment (Figure 2). Cisco IT will upgrade to 192 GB RAM to achieve higher virtualization densities.

Figure 2. High-Density Computing Conserves Floor Space in Cisco Data Center – Raleigh 1



The Cisco UCS chassis also contains UCS B440 Blade Servers for database environments that have high transaction-arrival rates, including mission-critical Oracle eBusiness applications.

Cisco IT created multiple Cisco UCS service profiles for each type of server used in Cisco Data Center – Raleigh 1:

- Development
- Test
- Staging
- Disaster recovery VMs for Oracle eBusiness applications
- Disaster recovery physical servers for Oracle Real Application Clusters (RAC)

Server administrators can apply a service profile to any combination of blade servers with a few clicks. “Cisco UCS service profiles are one of the main enablers for the dual-purpose data center, giving us the flexibility to quickly configure servers for DR instead of non-production, and then revert,” says Kahlon. “In the event of a disaster, we can take down the non-production VMs to create more capacity for DR VMs. Making the cutover on the server instead of the network saves time.”

Cisco IT maintains 20 percent floating capacity on the Cisco Unified Computing System, both on physical Cisco UCS servers and VMware ESX servers. The capacity is used for DR testing as well as preprovisioned foundational services used by all business applications. “Floating capacity also enables us to migrate critical code to the DR production environments in the event of a regional disaster,” Kahlon says.

All Cisco UCS B200 M2 Blade Servers connect to the storage area network (SAN) as well as network-attached storage (NAS) through a single pair of Cisco UCS 6120 Fabric Interconnects, and UCS 440 Blade Servers connect through Cisco UCS 6140 Fabric Interconnects. The fabric interconnects connect to the SAN by way of the Cisco MDS 9513 Multilayer Director, over four 8Gbps Fibre Channel interfaces. The fabric interconnects connect to NAS through a Cisco Nexus 5000 Switch.

Cisco UCS blade servers connect to the fabric interconnects through a Cisco UCS M81KR Virtual Interface Card (VIC), which is a converged network adapter (CNA). Cisco IT can configure each VIC with 128 interfaces, in any combination of virtual Ethernet interfaces or virtual host bus adapters (HBA) needed to meet application requirements. When using Cisco UCS M81KR Virtual Interface Card, Cisco IT took care to understand the relationship of the virtual HBA to the physical server. “If you have eight virtual HBAs, each of two LUNs [logical unit numbers] might be using four,” says Nunn. “We made sure that all four for a particular LUN were not tied to the same processor.”

Cisco Nexus Architecture: Non-Production and DR Environments Share Same Physical Switches

Cisco Data Center – Raleigh 1 has a 10 Gigabit Ethernet end-to-end network, built from Cisco Nexus 7000 and 5000 Series Switches, Nexus 2000 Fabric Extenders, and Nexus 1000V Virtual Switches.

Like Cisco UCS service profiles, the VRF feature of Nexus 7000 Switches is a critical enabler for repurposing the data center. VRF enables non-production and DR VLANs to remain operational at all times, without requiring separate Cisco Nexus 7000 Switches for each.

“Historically at Cisco IT, we built both an internal farm and DMZ farm for networking and storage,” says Mike Martinelli, data center engineer for Cisco IT. “VRF has enabled us to consolidate two farms into one, providing both sets of services.” All hardware and software related to VMware (physical ESX servers, data stores, and management software) are on the internal farm. Only the VMs for the DMZ environment are external-facing.

“When we are not in DR mode, the non-production VMs belong to VLANs in the non-production VRF segment,” says Wilson Ng, IT engineer for Cisco. “In the event of a regional disaster in the Texas MVDC, an automated process takes down the non-production VMs and brings up the DR VMs in the VLANs in the DR VRF segment. Cisco IT does not need to change access control lists, so we need only about 10 minutes to change the server personality from non-production to DR.”

Cisco Nexus 1000V Virtual Switch

A software switch embedded in the VMware hypervisor, the Cisco Nexus 1000V Switch provides the same capabilities as a physical switch, and also keeps networking policies associated with VMs as they move to different servers within the system.

“Traditional virtual switches have a per-ESX server configuration, meaning that Cisco IT would have to touch all 2000 VMware ESX servers in the event of a farm change,” says Martinelli. “The Cisco Nexus 1000V configuration applies to up to 500 ESX servers. That means configuring 500 ESX servers takes 10 minutes instead of 40 hours.”


Storage Area Network Design: Edge-Core-Edge

The SAN design objective was to minimize storage footprint to make best use of limited space and power. Raleigh 1 has 2.88 megawatts of power, compared to 5 megawatts in Cisco's production data centers.

Non-production and DR data coexist on the same set of arrays. “Using the same array for DR and non-production will enable us to use the continuously updated DR data as a source for development,” says William Nunn, data center design engineer for Cisco IT. “Maintaining both data stores on the same array will significantly reduce replication time to develop instances. Also, cloning multiple related instances that share common data may reduce physical storage requirements.”

Cisco IT used a standard edge-core-edge design for the SAN, with pairs of switches for hosts, gateways, and storage (Figure 3). A single SAN fabric in Cisco Data Center – Raleigh 1 supports four virtual SANs (VSANs): physical servers used for non-production, virtual servers used for non-production, physical servers used for DR, and virtual servers used for DR.

Figure 3. SAN Fabric



The Fabric Port (F-Port) Trunking feature of the Cisco UCS 6120 Fabric Interconnect enables a single F-Port to participate in two VSANs. F-Port Trunking requires dedicated rate mode, so Cisco IT uses 4-Gbps feeds on ports that face the Cisco Unified Computing System. For UCS B200 servers, Cisco IT aggregated all four ports to a single 16-Gbps trunk. Cisco UCS B440 servers connect through 32 Gbps or larger port channels.

“Without F-Port trunking, our design dictated that a Cisco UCS cluster could connect to only two VSANs per fabric,” says Justin Bruce, data center design engineer for Cisco IT. “F-Port Trunking lets us use more VSANs, giving administrators more flexibility in assigning individual blades to different VSANs. This feature also enabled us to evenly distribute storage traffic across available uplink ports, optimizing performance.”

Cisco IT also set up Cisco Fabric Services (CFS) regions to separate Cisco Data Center – Raleigh 1 device aliases to prevent them from transferring to other fabrics. “CFS regions enable us to hold the data locally and make sure we’re not pushing it out to Texas,” Bruce says.

Cisco IT currently monitors the storage environment using Cisco UCS Fabric Manager and NetApp SANscreen, and will later use Cisco Data Center Network Manager.

Physical Plant

Cisco Data Center – Raleigh 1 is a Tier 2 facility occupying 18,500 square feet (1719 square meters). Major design decisions included:

- **230V to equipment racks:** “Compared to the traditional 208/120 volt distribution, 415/230 volts to the racks is more efficient, costs less to operate, and eliminates the need for one level of transformers in the data hall,” says Mike Lavazza, engineering manager for Cisco Workplace Resources.
- **Chemical-free water-treatment process:** Cisco IT implemented a non-chemical alternative for the control of scale, corrosion, and biological bacteria in the condenser water system. A high frequency electronic pulse passes through the water, stripping suspended solids of their negative static charge. After the charge is removed, the remaining solids are precipitated as a powder and then filtered out.
- **Overhead cabling:** All data hall services, including the airflow system, power, and data cabling, are delivered from an overhead grid above the cabinets (Figure 4).

Figure 4. Cisco Data Center – Raleigh 1 Does Not Have Raised Floors, and Overhead Cabling Simplifies Changes




- **Waterside economizer:** Another way Cisco IT is reducing greenhouse gas emissions in Raleigh 1 is by using a waterside economizer. When air conditions are appropriate, the system uses a plate-and-frame heat exchanger instead of the chillers.
- **Multimode fiber infrastructure:** The data center supports 10 Gigabit Ethernet today, and can also support future 40 and 100 Gigabit Ethernet standards.

- **Physical access control:** Cisco Data Center – Raleigh 1 uses dual authentication for physical access via a card-based system and biometrics.

Results

Figure 5 summarizes the differences between the new Raleigh 1 facility for disaster recovery and the Cisco Data Center – Raleigh 5 facility it replaced.

Figure 5. Cisco Data Center – Raleigh 1 Provides More Automated Disaster Recovery Capabilities for the Texas MVDC



Rapid Disaster Recovery

Using automated processes, Cisco IT can recover all production applications in the Texas MVDC, using NetApp and EMC tools in conjunction with VMware Site Recovery Manager (SRM). “A push of a button starts the recovery process,” says Martinelli. Cisco IT uses Oracle Data Guard to replicate Oracle RAC databases.

Goals for disaster recovery vary by application priority, reflecting how critical they are to Cisco’s business:

- Priority 1: Active-active
- Priority 2: 4-hour Recovery Time Objective (RTO); 1-hour Recovery Point Objective (RPO)
- Priority 3: 24-hour RTO; 1-hour RPO

Efficient Resource Utilization

“Pre-provisioning capacity for DR in Raleigh 1 would consume resources and also increase operational expense because of constant patching,” says Kahlon. “Instead, we use VMware SRM to keep the source and destination in synch. We don’t use capacity until we make the decision to failover.” If a disaster occurs, Cisco IT initiates an automated process to suspend non-production resources, freeing them up for disaster recovery.

“Traditionally, companies have had to make a significant investment in DR and business continuity systems,” says Cribari. “The options were to build a DR site, or else contract with a service provider. By finding another option, designing a data center that could serve dual business purposes, Cisco IT avoided significant costs.”

Environmental Sustainability

Cisco IT is seeking the U.S. LEED-NC 2.2 Gold certification for construction, recognizing design for energy savings, water efficiency, CO₂ emissions reduction, improved indoor environmental quality, and stewardship of resources.

Among the ways the Cisco Data Center – Raleigh 1 site received points are:

- Power Usage Effectiveness (PUE) of 1.4 at full load
- Chimney rack hot-air isolation
- Waterside economizer, used 41 percent of the time
- Variable Frequency Drives (VFDs) on major equipment, including pumps, chillers, and computer-room air handler (CRAH) units
- Service distribution of 480Y/277V; rack distribution of 415/230V
- Photovoltaic cells on building roof
- Heat recovery from data hall for office space use
- LED exterior lighting
- Low-E glass windows
- Ability to use reclaimed water in cooling towers
- Non-chemical water treatment system
- Water-efficient plumbing
- Occupancy sensors integrated with lighting and temperature controls
- Landfill diversion during construction

Next Steps

Based on the performance of VRF in the non-production environment in Raleigh 1, Cisco IT will consider using it for multitenancy pods in production data centers. “The ability to virtualize our environments on pairs of Cisco Nexus switches saves the costs of separate switch pairs and minimizes space requirements,” says Cribari.

Another plan for Raleigh 1 is using Locator/ID Separation Protocol (LISP), a feature of the Cisco Nexus 7000 Switch, to optimize routing between the Texas MVDC and Raleigh 1. Cisco IT has already conducted a proof of concept using LISP for long-distance WAN VMware SRM recoveries over Metro Ethernet. “With vMotion, we recovered over 800 miles in just 6 minutes, due to not having to change IP addresses,” says Martinelli. “There was no interruption other than the actual outage.”

Finally, Cisco IT continues to move forward with the Cisco IT Elastic Infrastructure Services (CITEIS), Cisco’s private cloud, used for Infrastructure as a Service (IaaS). On the roadmap are Software as a Service (SaaS) and Platform as a Service (PaaS). Cisco Data Center – Raleigh 1 will use image-based deployment, part of CITEIS, for rapid deployment of non-production applications with complex configuration requirements.

Lessons Learned

Cisco IT shares the following lessons learned with IT teams at other organizations considering multipurpose data centers:

- Think in terms of superpods that support multiple environments. The enabling technology is the VRF feature of Cisco Nexus 7000 Switches.
- Double-check business and technical requirements to size the data center appropriately. For Cisco Data Center – Raleigh 1, the DR and non-production environments had distinctly different requirements.

- Do not underestimate the task. “Designing a multipurpose data center is not as easy as it sounds,” says Nunn. “We attribute our success not only to technology, but also to building and developing partnerships with our strategic stakeholders.”
- Build a DR solution using technical and operational competencies that you already have.
- Consider the business impact of shutting down non-production infrastructure to use it for disaster recovery. Cisco IT decided the cost benefits of a repurposed data center outweighed lost productivity for development engineers on the infrequent occasions their servers would be shut down.
- Thoroughly document processes for disaster recovery. “Operations engineers need detailed procedures to change the data center from one purpose, such as non-production, to another, such as disaster recovery,” says Cribari. The procedure should include the sequence of applications to recover.

For More Information

To read additional Cisco IT case studies on a variety of business solutions, visit Cisco on Cisco: Inside Cisco IT www.cisco.com/go/ciscoit.

To learn more about the Cisco Unified Computing System, visit www.cisco.com/go/ucs.

Note

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described; Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties, therefore this disclaimer may not apply to you.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)