



WHITE PAPER

DATA CENTERS: BEST PRACTICES FOR SECURITY AND PERFORMANCE

EXECUTIVE SUMMARY

Network service providers seeking Cisco® Powered Network Program membership must meet rigorous application requirements that, among other things, include network performance, network security, and customer support. Designation also requires an annual program compliance review that guides the Cisco Powered Network Program member toward best practices in the area of overall network design and operations. In the course of working with service providers around the world on these technical reviews, Cisco has developed a criteria and checklists to determine program eligibility. This Data Center White Paper is useful to those constructing or managing data centers with stringent performance and service requirements. Enterprise and commercial customers seeking to purchase Web hosting, applications hosting, call and contact management, or content delivery-type services may also find this guide helpful in evaluating a service provider's data centers.

DATA CENTER DESIGN OBJECTIVES

There are four objectives in the design of any high performance data center: Security, availability, scalability, and manageability.

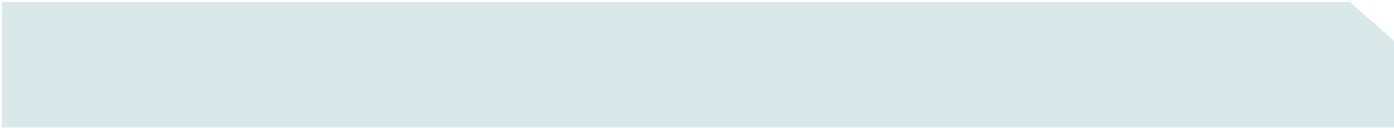
Today's business and competitive environment requires that each of these objectives be considered from a complete end-to-end perspective. Clients or users connecting to the data center measure performance by timely access to the desired application data, whether the connection is point to point or via an Internet connection. The user must perceive reasonable response and connection time. Access is no longer measured solely on an arbitrary measure of "network availability" or ping times. This fact does not lessen the goal of developing networks that provide 99.999 percent availability. However, the implication is that the designers of the traditional networking elements must also become "content-aware."

Security

Security can no longer be treated as an afterthought when it comes to optimum network design. Trial-and-error networking is not an option, as a single vulnerability could compromise the corporate lifeblood, the network. In such a high-stakes environment, only specialized expertise is acceptable. There are challenges in offering encryption, certification, directory, network, and other security components that enable what one would consider a 100 percent secure network. While industry struggles with developing the technology to provide these protective components, the IT manager must still cope on a daily basis to reduce the network's imminent risk.

A complete network security solution includes authentication, authorization, data privacy, and perimeter security. Perimeter security is traditionally provided by a firewall, which inspects packets and sessions to determine if they should be transmitted or dropped. In effect, firewalls have become a single point of network access where traffic can be analyzed and controlled according to parameters such as application, address, and user for both incoming traffic from remote users, and outgoing traffic to the Internet.

In general, firewalls are intended to protect resources from several kinds of attacks such as passive eavesdropping/packet sniffing, IP address spoofing, port scans, denial-of-service (DoS) attacks such as SYNchronize-ACKnowledge attack (SYN flooding), packet injection, and application-layer attack.



See the data center best practices checklist (Appendix A). The checklist begins with the physical security of the data center and its environment (electric power, temperature, and humidity controls). The checklist also contains recommendations concerning network security in general terms, recognizing that threats may occur actively or passively (sometimes called opportunistically), and from people and events either outside or inside the data center itself. Connectivity paths to external networks, and each function inside the data center (such as gateway or WAN edge, core, distribution, and access), must be designed with no single point of entry for unauthorized users.

Availability

Availability is generally ensured by the overall network design and implemented in several ways. First, networks are designed with steps to minimize the occurrence of service problems and the time to recover from problems (such as backup recovery policies). Second, high availability must be considered at each layer of the Open Systems interconnection (OSI) reference model, with redundancy and failover provisions made at the physical, data link (for example, Ethernet), network (for example, IP), and application layers. The most effective solutions are those with consistent engineering considerations tightly integrated throughout the data center, rather than those approached with a series of “point-solution” products and techniques.

Scalability

Scalability must be provided in every data center. Server load balancing is the norm, and techniques such as reverse proxy caching are often used to offload servers. Server load balancing, like other aspects of data center design, must also be content-aware, preferably using delayed binding, full URL and cookie inspection, and “sticky (server) connections” as part of the logic of choosing a server for each user request. Specific content may be in high demand and considered “hot.” Because this content may not be known in advance (such as breaking news stories), advanced data centers should also have the ability to automatically and immediately identify and replicate hot content to overflow or backup servers or cache to ensure the ability to support the increased demand without compromising performance.

Both the complications and available solutions increase greatly if the data center is part of a geographically dispersed set. The ability to provide content at multiple sites allows the served data to be “closer” to the requesting client, thus providing faster response but also higher availability (due to partial or complete redundancy of the data). But multidata center design must then include considerations of management of the data content, distribution of updates, synchronization of different sources, proper routing of requests, handling of downed servers, additional security and so on.

Manageability

Manageability means much more than simply knowing if a server or other network element is “up or down.” Especially in service provider data centers supporting multiple customers, the ability to assess service levels on a per-customer basis is essential to the offering and administration of service-level agreements (SLAs). Managing IP address assignments, keeping track of network configurations, and not losing site of trouble tickets and alarms often requires use of a mechanized network operations center (NOC) support system. Good manageability tools and qualified personnel supporting the infrastructure translate into lower operations costs, since time is not wasted trying to resolve indications from conflicting management systems and higher customer satisfaction.

APPENDIX A

Cisco Powered Network Program Data Center Best Practices Checklist

Facility and Physical Requirements

- Multiple physically separate connections to public power grid substations
- Continuous power supply with backup uninterruptible power supply (UPS) systems:
 - Adequate UPS capacity including air conditioning and lights
 - UPS systems tested at full load on monthly schedule
 - Fuel for UPS generators (48 hours worth) kept on premises and monitored for local environmental compliance
- Conform to or exceed applicable local structural building codes utilizing standards such as bullet proof glass, fire doors and reinforced walls and complying with disaster proof design:
 - Comply with all local zoning ordinances
 - Certify not located in a 100-year flood plain
 - Earthquake and hurricane bracing on all racks and cable trays (where appropriate)
- Adequate multizone air conditioning, including a backup system for the multizone air conditioning:
 - Climate control including humidity sensors and control
- Heat and smoke detectors that meet or exceed all local fire code regulations.
 - Very Early Smoke Detection Alarm (VESDA)
 - FM200 [ETG5] fire suppression system in data center and NOC
 - Separate detection/FM200 zone under raised floors
- Preaction dry pipe system zoned to release water only where needed
- Easily removable access panels in raised flooring
- Flood sensors and monitoring under raised floors and in other critical areas
- Separate grounding systems to prevent grounding loops; true ground versus green wire ground
- Sealed cable vault entrances to facility, remotely monitored
- Formalized physical facility preventive maintenance program
- Sub-breakers per relay rack or lineup
- 48 VDC power converters, 220 VAC, 20A, 30A, 40A
- Power filtering in UPS system

Physical Security

- Written security policies readily accessible:
 - Badge sharing and piggy back entry rules
 - All visitors must be admitted through reception
 - Written statement of work upon sign-in

- Building access procedures:
 - Limited number of building entrances in compliance with local fire ordinance
 - Provide access to limited and managed security policies for all facility entrances
 - 24x7 onsite security guards
 - Visitor-logging procedure
 - Card-key, biometric, or similar entry locks
 - ID-badge system for all employees and visitors
 - Staff and visitors must wear badges at all times on premises
 - Individual cabinet locks; master in NOC; key list from customer
- Equipment locations:
 - Video surveillance and motion sensors for entrances, interior doors, equipment cages, and critical equipment locations within the building
 - Locked cages with ceilings; locking cabinets with climate control for those wanting more privacy
 - Secure rooms available
 - Managed firewall services with 24x7 monitoring available
 - Backup lighting systems for entry ways and cable vaults
 - Individual cabinet locks; master in NOC; key list from customer

Network Security

- Written network access security policies readily accessible:
 - Password policies (such as not sharing, lengths, forced renewal, aging)
 - Acceptable use (ISP not allowed to run programs that are illicit or illegal; use of sniffers or cracking/hacking programs are not required)
 - Documented user responsibilities on security in company policies and re-enforced by education
 - Asset protection
- Network security infrastructure in place:
 - Perimeter protection (firewalls, filtering router)
 - Intrusion detection
 - Authentication and authorization (passwords, RADIUS/TACACS, Secure IDs)
 - Backup and recovery systems to restore after a problem, such as load balancing, failover protection
 - Regular assessment of network infrastructure
 - Assessment of network expansions or additions
 - Tape or media storage offsite backup
 - Regularly scheduled security audits
 - Server antivirus software protection

Operations

- Database of all installed equipment and configurations
- Toll-free telephone support

- Supported monitoring:
 - 24x7 monitoring of dedicated servers and network equipment (note both frequency and method, such as PING, Simple Network Management Protocol [SNMP])
 - 24x7 monitoring of the health of the equipment with alarms and pager alerts for network failure and failovers
 - 24x7 monitoring firewall services available
 - Alternate NOC available
 - Second-tier support personnel located nearby
- Trouble ticket processes:
 - Created and logged for all unusual or unexpected events
- Automated case escalation procedures in place including escalation timeframes
- Reporting that provides trending statistics on trouble tickets and minutes (above) to facilitate quality and customer reports
- Performance reporting and end-user impact monitoring
- Periodic and exception reports provided to customers (including usage and problem reports)
- Spare equipment on site for key networking equipment available in case of hardware failure
- Business continuity plan:
 - Daily site backups
 - Tape vaults or other secure storage facilities on site in case of natural disaster
 - Onsite and offsite storage available
- Customer callout and escalation database
- Intercom system
- Written procedures for each customer on alarm handling

Backbone Connectivity

- Multiple direct connections to Tier 1 Internet carriers using high speed Cisco routers as gateways
- Border Gateway Protocol vs. 4 BGP-4 routing
- Class C Internet address blocks available
- All backbone services have Cisco Powered Network Program member designation
- Each carrier has a secure termination area, and location supported via the NOC or the carrier providing the termination
- Fiber enters the data center through diverse conduits or routes (for example, if a backhoe cuts through conduit, the network reroutes to minimize loss of service)
- Aggregate bandwidth sufficient to scale the network to meet customer's service demands
- Describe policy on facility utilization or over-subscription
- Provider must have private facilities connecting to other data centers, and a documented process
- Multiple Internet access
- Carrier XCONN and distribution system; separate carrier point-of-presence (POP) area
- Formalized SLA policies

- Roof rights and riser conduit right of way
- Multiple riser conduit from cable vault to data center

Gateway/WAN Edge Layer

- High-end routers (such as Cisco 7500 or 12000 Series) in a redundant configuration
- Cisco Hot Standby Router Protocol (HSRP) implemented
- BGP-4 implemented
- Adequate total packet-per-second capacity for peak customer load
- Firewalls in place
- Network security team in place
- Remote firewall management offered

Core Layer

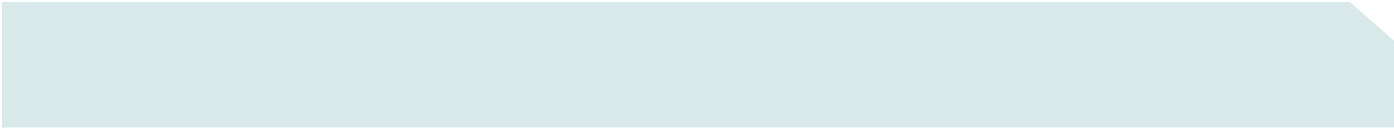
- High-end switches (such as Cisco Catalyst® 8500 or 6500 Series) deployed
- Switching and links entirely redundant with no single points or paths of failure
- Web cache redirection implemented
- Content and Transmission Control Protocol (TCP) offloading implemented via reverse proxy caching
- HSRP implemented for fail-over protection
- Intrusion detection implemented (such as Cisco Intrusion Detection System)
- Automatic notification of intrusion attempts in place

Distribution Layer

- High to mid-range switches (such as Cisco Catalyst 6500 or 6000 Series) deployed
- Switching and links entirely redundant with no single points or paths of failure
- Caching systems (such as Cisco 500 Content Engine or 7300 Series) implemented
- Server load balance (Cisco CSS 11000 Series) implemented
- Server content routing (Cisco 4400 series) implemented if multiple data centers
- Caching (Cisco Cache Engine 500 or 7300 series) implemented

Access Layer

- Mid range switches (Cisco Catalyst 6000 or 4000 Series) deployed
- All servers dual homed



Cabling

- All cable runs located under raised flooring and appropriately marked
- All cable runs physically protected from damage via tie-downs or where appropriate in conduit
- All cabling designed to Category 6 specifications (to support 1-Gbps data rates)
- Communications cabling raceways separate from electrical; no intersections
- Shielded cabling for T1/T3s. DSX panels for XCONN, demarcation, and test points
- All cabling on raceways, tied down

**Corporate Headquarters**

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco Website at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica
Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR
Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia
Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan
Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2004 Cisco Systems, Inc. All rights reserved. Catalyst, Cisco, Cisco Systems, and the Cisco Systems logo are registered trademarks or trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0406R) SB/LW6660 07/04