WHITE PAPER

CDW·G PEOPLE WHO GET IT™

CONTINUITY OF OPERATIONS

Virtualization, agile development, increased mobility and cloud computing are challenging government agencies to extend COOP beyond disaster recovery concerns.

Executive Summary

Not since the rise of the Internet in the 1990s has there been such dramatic change in the IT industry. Cloud computing, virtualization, mobile computing and a new orientation to applications as services have all combined to change the way government agencies ensure continuity of operations (COOP), high availability and peak performance of their applications and networks.

Not long ago, disaster recovery (DR) and COOP were synonymous. The IT staff could confidently provision backup servers at alternate data center sites and install failover power supplies, and it would be set. While backup and recovery facilities for DR certainly remain foundation blocks of COOP planning, changes in system architecture have pushed IT staffs to rethink their strategies and solutions.

Today, virtual machines move within data centers and across wide area networks (WANs). Agile development and a move toward a simplified approach to applications mean that logically related resources are often scattered physically.

TWEET THIS!

Table of Contents

More people also are working remotely and on the go, using a variety of devices and wireless services. In addition, while cloud computing increases flexibility and scalability, it also makes organizations more dependent on links to external facilities.

These developments take COOP well beyond the boundaries of disaster recovery. Continuity planning requires optimizing WANs, guarding against single points of failure in complex network topologies, and ensuring that users receive high and reliable levels of service. In today's world, IT staffs must tightly integrate COOP into their plans to maintain robust availability and still boost efficiency.

Data Access and Storage

Government organizations often take a document-centric approach to their missions. For example, applications for benefits, insurance and healthcare all require forms, both online and paper. Within many agencies, case management systems are still largely paper-intensive. In the course of facilitating field reports, taxes, fees, fines, licensing and permitting, and land and vehicle registrations, government agencies generate billions of documents.

Recent studies have found that documents and unstructured data are growing 30 percent per year among large organizations. These documents exist in three forms: paper, unstructured electronic (such as PDFs) and structured electronic (such as those found in databases). Regardless of what form its documents take, an agency needs them readily available to conduct the daily activities of government. And it must also retain documents for archival and legal-compliance purposes.

These concerns make secure capture and storage of documents as well as access to the data they comprise crucial to continuity of operations. Agencies rely on document capture for the DR component of COOP because the process converts paper documents into digital versions they can then back up along with their other data.

IT staffs also depend on document management to maintain the high availability and application performance components of COOP. Document management gives organizations a way to tag and categorize data contained in unstructured electronic documents so that the needed information is more quickly retrievable.

Agency IT departments can take any of several routes to incorporating data storage into their COOP plans. Disk mirroring (simultaneously writing data to two physical disks) and redundant arrays of independent disks (RAID) are the primary technologies used for ensuring continuous availability of production data. Others approaches include the following.

Data tiering: Also known as hierarchical data storage, this strategy prioritizes data according to age or frequency of access. Production data sits at the highest tier, meaning storage on the fastest and generally most expensive level (typically solid-state disks or fast primary magnetic disks).

Depending on the array, this tier level can be broken down into several individual tiers, based on I/O performance. Such a breakdown might include solid state drives (SSDs), 15,000 rpm serial-attached SCSI (SAS), 10,000 rpm SAS and 7,500 rpm SAS hard drives.

As data moves from production to stored availability to archiving, the material moves physically to slower and less expensive media (generally, tape and optical). As a COOP matter, IT staffs may store archival data off of the production site.

Data deduplication: A form of file system management, data deduplication improves application performance and reduces bandwidth usage, thereby increasing network availability. The basic concept behind deduplication is that it is inefficient to write multiple copies of the same data to storage.

A prime example of this would be a Microsoft Office document that's in the possession of two or more users. The IT staff only needs to write one copy of the file to storage or backup, with users having pointers to that single instance of the document. Several deduplication techniques exist. The primary differences between them depend on whether they operate

IaaS, SaaS and Continuity of Operations

When a provider packages data center infrastructure and all the elements it entails and offers it via the Internet, the result is infrastructure as a service (IaaS). More organizations are using IaaS providers as a way to gain flexibility for scaling capacity up or down within a predictable cost model.

If that sounds a lot like cloud computing, that's because IaaS is a basic building block of clouds. IaaS can also serve as a shortcut to a DR site as part of a COOP plan. The IaaS provider is a separate entity devoted to a high level of uptime in a competitive market, and so it's unlikely that a disaster that strikes the organization would also affect the provider. It's essential to thoroughly examine the IaaS provider's COOP plan before signing on.

The larger providers typically operate several geographically separate sites, both to minimize WAN distances and to isolate whatever problems might occur at any one of them.

Adding an enterprise application to IaaS makes it software as a service, or SaaS. The SaaS industry has its roots in applications built specifically for remote hosting. Customers access the application and their own associated data files via the Internet, paying per user per month. Today, cloud providers often host backup instances of customers' own applications, delivering an instant COOP option.

at the block or file level and inline or post-process (after the data's storage). Unstructured data and backup files are ideal candidates for deduplication.

Storage virtualization: This technique pools all available storage as a single resource that the IT staff can manage centrally. Storage virtualization inserts an abstraction layer that supports capacity and performance management, as well as data protection services, without the need to deal with the complexities of physical storage and multiple providers.

High Availability and Failover

Most government agencies have moved firmly into the Web 2.0 era. That means they conduct more citizen engagement online, whether that means gathering input on a new regulation or conducting transactions such as tax payments or fee collections via the Internet.

So high availability — including interruption-free failover for when something does crash — remains at the top of the list for an IT staff's data center operational goals.

In theory, every application could have its own server and backup, coupled to redundant network channels. This would create a bulletproof system, but also one that few agencies have the budget for.

Every program office or division within an agency can find reasons to justify having its own data center. And for years, many did. The financial crisis of 2008 changed all that. In today's environment, with government at all levels facing tough fiscal conditions, data center consolidation has emerged as a primary IT cost-cutting strategy.

At first glance, collapsing two or more data centers into one would appear to work against disaster recovery and COOP by creating fewer, but more critical, points of failure. On the contrary, data center consolidation offers an opportunity to update or reengineer an agency's approach to availability and disaster resistance.

Consolidation brings together facilities that were separated by distance. It primarily improves availability and disaster protection by reducing complexity. Fewer data centers require fewer WAN links, LANs and storage subsystems.

Data center consolidation also provides a logical point for updating operating systems, backup tools and failover software. Many data center operators are anticipating the arrival of a new generation of servers equipped with Advanced RISC Machine (ARM) chips — the reduced instruction set computer architecture used in mobile devices. Operating at lower temperatures and power inputs compared to x86 hardware, ARM technology promises to increase availability by allowing for simplified cooling systems and lighter power requirements in the data center.

Consolidation results more in the pruning of small, underutilized, obsolete data centers, not the merger of all centers into a single one. Many organizations also use consolidation to weed out unneeded, redundant or obsolete applications that present security vulnerabilities and take up resources. If nothing else, these cleanups aid availability by lightening server backup and recovery processes.

5-step Approach to COOP

Action	Purpose
Prepare	Adopt early-warning tools and continuity and situational response plans
Prevent	Safeguard staff, citizens, property and assets
Detect	Provide instant notification of disruptions, security breaches and threats
Assess	Determine the scope of the incident and the next actions
Respond	Coordinate real-time communication

SOURCE: Cisco Systems

Most IT managers will agree that more than any other trend, server virtualization makes consolidation possible. On the other hand, the consolidation mandates issued at all levels of government trigger many agencies to move ahead with virtualization.

Here again, reducing physical assets improves an organization's COOP plan. Consolidation coupled with virtualization can also become the path to an organization's growth when existing, nonvirtualized data centers reach their power capacity and physical-space limits, a point at which reliability and availability can suffer.

Server virtualization produces high availability in several ways. Portability tops the list. Virtual machines (VMs), typically applications and their associated operating system (OS) and memory resources, exist as pure software objects. The relatively small amount of hardware resources individual VMs require, plus their self-contained quality, allow IT managers to easily replicate them as backups. These features also allow IT staff to move VMs from machine to machine.

Virtualization providers include various management tools, products that are also available from third-party utility publishers. This software lets administrators move VMs among servers within a cluster, servers on a LAN within a data center and even among data centers across the WAN.

VM portability supports server load balancing, and VM replication coupled with portability makes for instant availability of a backup machine in the event of a VM failure. The cause of the failure doesn't strictly matter. At times it could be from a software error or corruption of some sort, or possibly a hardware crash.

Virtual machines permit a fine-grained approach to restoration. Administrators may bring back a VM at the file or application level. They may simply restart the entire VM. (Keep in mind, this is a faster and more efficient procedure than a hardware restart.)

Or they may configure backup and restore functionality to go live with a backup copy. What this all means is that availability procedures can match the situation. Not every backup call is for DR. In critical online environments, maintaining response times and application service levels requires that kind of restoration flexibility.

Maintaining high availability also depends largely on choosing the right backup tools. Utilities designed for full-server backups can by definition back up virtual machines. It's important to see if they are rooted in the physical backup world. If so, they might not have the deep VM-specific functionality the IT staff needs. The leading virtualization providers include application programming interfaces that allow utility manufacturers to write special backup features.

It's worth evaluating the many backup tools available. The most capable ones will include image-based backup, file-level restoration and deduplication. The last feature reduces restoration times and preserves network bandwidth.

Some packages let IT managers run applications directly from backup machines for zero data loss failover. Published case histories report failover times as short as one VM ping — down from 30 seconds (which can be an eternity in production and mission-critical environments).

Also look for functionality in backup software that checks the "heartbeat" of backed-up virtual machines to verify their functionality and true standby readiness.

Data center managers have always considered hardware independence a virtue. When applications can run on any hardware, organizations are free to seek bids for the most inexpensive hardware sufficient to host their software.

Chip advances now make full virtualization of 64-bit applications possible, bringing AMD and Intel cross-compatibility to these apps. Primarily a management advantage, this compatibility also allows more flexibility in designating sites available for backup and recovery.

Networking

Ubiquitous traffic cameras on the roads and highways let authorities know exactly where bottlenecks originate. They can dispatch resources such as tow trucks more quickly when they can see what's going on. Similarly, if IT managers can see clearly what's going on in enterprise networks, they get early warnings about traffic overloads, device failures and threats from rogue user software or malware.

That's where remote monitoring (RMON) comes in handy. RMON tools let system administrators view network dynamics as they play out. This visibility contributes much to high network availability and COOP by letting organizations fine-tune their risk management approaches. In fact, rather than run periodic security audits, federal policy now states that agency technology shops should monitor their networks continuously.

A denial of service (DoS) attack ranks high on the list of events that can slow down or stop an agency's operations, so ensuring a strong security posture contributes greatly to COOP. In recent months, vulnerabilities in virtualization hypervisors have emerged, and the influx of mobile devices at agencies has engendered new strains of cyberattacks. Phishing via Skype, robo calls and texting have been added to the arsenal of virtual weapons employed by cybercriminals.

Remote, continuous monitoring also gives network administrators the ability to head off other problems, including malfunctioning devices, failed applications, policy violations and network bottlenecks.

Monitoring basically looks at two categories of network information: packet traffic flow and devices. Tools for monitoring may operate using a dedicated appliance or an application running on an existing switch. Either way, RMON should yield information about device performance, traffic between any two points on the network and anomalies existing within traffic.

For finding anomalies in traffic, deep packet inspection is an aspect of RMON that's growing in importance. Traffic flow gives IT managers the basics of bandwidth constraints as well as some packet source information. A new crop of products now looks inside the packets, either in real time or cached as a data set.

Deep packet inspection means more work for the network staff, who must analyze the results of the inspection. It also requires dedicated storage for the cached packets. A typical cache contains at least two days' worth of traffic, so plan on adding about 2 terabytes of storage to support deep packet inspection.

Two more important COOP concerns are load balancing and redundancy. The following products and tools help IT managers to stabilize the network.

Load Balancing

Maintaining an acceptable level of network performance requires load balancing, a technique that keeps processor demand evenly distributed across available servers. Load balancing is important in both physical and virtualized environments.

Networks go down either on purpose (such as for maintenance) or by accident (because of a server crash).

With failover mechanisms in place, such events don't interrupt service. However, without load balancing, they can disrupt performance if a particular server sends all of its processing on a single server, instead of many others.

Load-balancing appliances typically come in a rack form factor. The load balancing generally occurs at Layer 2, 4 or 7 in the ISO protocol stack. Layer 2 load-balancing bundles link together to create more backbone bandwidth. Layer 4 load balancing distributes requests at the IP transport layer and is used to parcel traffic loads generated by web and other IP services. Layer 7 load balancing helps maintain quality of service for applications.

The load balancers divide processor-intensive work across a cluster of machines for high performance. Although these products have roots in scientific computing, the advent of big data analysis in other domains has made load balancing important to many organizations' daily activities.

Regardless of the application, any COOP and high-availability plan would be incomplete without a load-balancing component, especially if maintaining service levels is part of the plan.

Redundancy

Most IT managers consider redundancy a basic condition for COOP because no network component operates with the certainty of 100 percent uptime. The standard for annual uptime is 99.999 percent, or "five nines" (see *The Numbers on Five Nines* table). It's a tough standard to meet, allowing only about five minutes per year of outage — or about six seconds per week. But it is possible.

Nines calculations don't include scheduled downtime, but if IT managers plan outages to happen with no interruption in service, then they must build redundancy into the agency network. Redundancy requires planning and design, not merely doubling every device and link. It starts with analysis of which interfaces and applications are most important to mission delivery, then prioritizing the rest.

Redundancy and load balancing can take place simultaneously. For example, when an application runs on multiple servers, IT managers can set service levels higher across the enterprise and also provide for failover capacity should a server or application go down.

This is also where WAN optimization comes into play. To reduce network latency, application optimization requires IT managers to cache frequently used data near endpoints, or at least near LANs encompassing endpoints.

And to optimize the network, IT managers must ensure that network traffic has multiple pathways between endpoints — again, to avoid congestion, while also providing redundancy. Together, redundancy, optimization and load balancing add up to resiliency, the quality that IT managers want in a highly available network that can recover from disasters.

The Numbers on Five Nines

Availability Level	Downtime per Year
99.999%	5.25 minutes
99.99%	52 minutes
99.9%	8.76 hours
99%	88.76 hours

SOURCE: Information Technology Intelligence Corp.

The growth in mobile device use also affects COOP planning. Given the growing number of mobile workers using Wi-Fi and carrier-provided data plans, it's important not to overlook wireless devices and services when planning for redundancy.

Bear in mind that Wi-Fi power comes from a local source, unlike cellular service. This means that redundancy for COOP and continued availability might require reliable cellular bandwidth within a building or complex if a power outage takes down the Wi-Fi network. Cellular repeaters require Federal Aviation Administration licensing. Enterprise-class repeaters may also require installation by specialized contractors.

At some point in the networking setup, the organization will deal with telecommunications carriers. Services from carriers

How Packet Inspection Aids COOP

Continuity of operations starts with preventing as many potential interruptions as possible. That's why so many government agencies are adopting continuous monitoring of their networks. Monitoring traffic flow is much like watching vehicular traffic on real roads. It can help the network administrator alter traffic patterns when bottlenecks develop.

Although continuous monitoring helps the cause, IT departments require more today. Deep packet inspection tools look inside and also route packets. For example, the tool may know the profile of a particular threat and block its passage through the enterprise firewall. Or a tool may simply record content in logs or cache the packets themselves in a database for later inspection.

The real goal for a COOP plan is not inspection itself, but rather situational awareness in as close to real time as possible. Deep packet inspection alerts the network staff to any of several conditions that could affect performance, including DoS attacks and loss of information.

Deep packet inspection can tell network sleuths what users were doing and when, such as who accessed what files, downloaded which app or removed data sets. In short, deep packet inspection offers a powerful tool to minimize insider threats caused by careless or malicious authorized users.

include long distance, high-bandwidth wire and wireless links to managed IP services within a building or campus.

Rather than pay for redundant links, it may make more sense to write strict service-level agreements and pay for the effect of redundancy. However, in all setups, the IT staff must incorporate a thorough testing plan. Testing should include physically disconnecting critical links and monitoring the resulting loads on the alternate pathways.

Client Access

The telework trend in the public sector has advanced quite a bit in the last few years. Several factors came together to support staff working outside of an agency's four walls.

To start with, the Telework Enhancement Act of 2010 gave a strong endorsement at the federal level, along with administration policymaking to implement the law. Several states have developed teleworking policies for their staff, and there's even a group called the Telework Coalition supporting state and local agencies. The Telework Exchange does the same for the federal community.

Although the drive to reduce carbon emissions has boosted support for telework, COOP is still its primary driver. Once workers are not required to work in a specific physical location to accomplish their tasks, the agency can remain operational even when its staff cannot access headquarters because of inclement weather or some other emergency.

On the technical side, the steady improvement in the speeds and ubiquity of broadband service, in addition to the advent of appealing endpoint devices, have transformed telework by blending it with the mobility trend. Telework used to mean working from home or from a telework center with presumably easier commuter access than offices in a central location. Today, it means working from anywhere.

A viable COOP plan must consider enterprise applications for mobile and remote workers. A growing number of organizations set up mobile devices, whether tablets or full-fledged computers, so that they don't store data locally.

However, when a department's data center resources are unavailable, it's often necessary to keep data in store-and-hold mode locally until there's an opportunity for remote workers to transmit the information. As an alternative, to better support cybersecurity strategies, remote users should have virtual private network (VPN) access to secondary data centers so they can work until the primary servers are back online.

Remote Access Options

With mobility becoming more common at many agencies, IT staffs are developing bring-your-own-device (BYOD) policies. In a practical sense, BYOD means IT infrastructure must support two basic device setups:

- Traditional machines with Intel or AMD x86 processors and hard drives that run Windows, Linux or OS X
- Tablets with ARM or similar reduced instruction set processors and small solid-state drives that run iOS or Android

Mobility affects application architecture, which in turn influences COOP planning. Client-server architecture is designed around the use of standard PCs or notebooks as remote clients and puts much of the application processing on the endpoints. In some ways, that setup made for more robust COOP because workers could do substantial work offline.

Enterprise adoption of the new mobile devices often becomes the trigger for moving on plans to virtualize users in the data center or cloud, and it gives users themselves thin-client access. The tablet screen presents a rendering of the application interface while all the processing occurs remotely.

This approach supports COOP in a new way. Most continuity planning focuses on the network and data centers. On the other hand, loss or theft of portable devices can also pose continuity problems.

These range from a single user's loss of productivity to the loss of sensitive or classified information stored on the device. IT managers also worry about the potential for a major network intrusion if a device that has been improperly secured falls into the wrong hands.

Secure Remote Access

The VPN has become the de facto standard route for remote users accessing networks. Carriers, equipment suppliers and third-party software providers all sell a variety of VPN products and services.

In its most basic form, a VPN is an encrypted channel that uses the Internet to connect a single user to the agency's network. Network administrators can tailor VPNs to each user's needs and access rights.

VPNs support COOP by virtue of their security. Data is encrypted en route. Coupled with two-factor user authentication, IT managers back this approach to make it less likely for an unauthorized person to gain access with a given device.

Cloud computing and client virtualization have made software as service (SaaS) increasingly popular. Users, in essence, get the functionality and benefits of software via an Internet connection. The IT department benefits because the software itself (and in the case of thin-client users, the processing) remains centrally hosted.

In some ways, third-party application hosting combined with browser or thin-client access creates the ultimate operational continuity. Even if an agency's offices are shut, locked and off the grid, the SaaS provider can provide service. IT departments that opt for SaaS should make sure the provider maintains

Are VXLAN and LISP Ready for Prime Time?

For IT managers, two primary issues raise many questions around COOP, high availability and productivity: facilitating reliable mobility of virtual machines (VMs) from server to server over wide area networks as a load-balancing and disaster recovery technique; and ensuring seamless, uninterrupted access as remote and mobile workers move from one wireless network to another.

A couple of new standards aim to meet these challenges.

A proposed standard for VM mobility was recently submitted to the Internet Engineering Task Force (IETF). The Virtual Extensible Local Area Network (VXLAN), proposed by Cisco Systems and VMware, lets VMs move over networks at Layer 2 (the data link layer) but still benefit from the services available from Layer 3 (the network layer). The standard was designed to move data across a network so that a VM still runs when it reaches its new destination, typically another data center.

VMware supports the VXLAN technology in vSphere 5 and vCloud Director 1.5.1. Cisco also supports VXLAN in its Nexus 1000V Series switches and the 1010-X Virtual Services Appliance.

Another solution makes delivering high availability to mobile users possible. Over the past several years, the Locator/ID Separation Protocol (LISP) developed by Cisco Systems has become a part of carrier and enterprise network setups to enhance the routing function on the Internet. LISP is also available to the IETF for continued open development. And at a recent conference, Cisco demonstrated how the protocol lets users stay online as they move out of reach of a Wi-Fi zone and into a wireless-broadband cellular network.

LISP also allows interruption-free mobile access to virtual machines as they move from a data center to a backup or branch site. This means mobile workers using cellular service or situated in a Wi-Fi hotspot could work continuously during a failover event.

escrow copies of applications and data files at an independent third-party site or at a secondary site of its own.

Unified Communications

At its essence, continuity of operations aims to maintain communications between workers and computing resources, the two basic elements in a knowledge organization, such as a government agency. Today, unified communications (UC), which encompasses both technology and applications, plays an increasingly important role in strengthening an agency's COOP plan.

If executed properly, a UC strategy will keep staff in touch with one another and with their applications during a disaster or emergency. And it will help ensure delivery of services by establishing robust and redundant data channels among endpoints and data centers. In short, any UC planning should take COOP into consideration.

UC consists of many components, including:

- **Communications infrastructure:** The IP network that carries voice, data and video traffic
- **Collaboration applications**: The apps that ride on the IP services, including voice and video conferencing, e-mail, text messaging, and instant messaging
- **Management tools:** The tools for controlling access to and provisioning of applications

"Unified" doesn't refer to complete physical convergence; organizations still need redundancy to fulfill COOP requirements. Rather, UC is a logical integration of services. For example, the multiple forms of communication available to a user are often imbedded into productivity applications, turning them into collaborative applications.

The Defense Department includes the term "highly available" in its definition of unified communications. And that's a good model for all government organizations that want to make UC a component in their continuity of operations and disaster recovery plans.

UC should also buttress the services an agency makes available to the public online. Disaster recovery that protects the internal processes of an agency during an emergency won't mean much if that agency has citizen-facing or interagency processes that become unavailable in a crisis. This is where the communication between computing resources aspect of UC matters.

For example, the Transportation Security Administration makes data available to airports and local transit agencies throughout the country. COOP and DR strategies for the TSA should cover both the internal and customer-facing components of operations. That's only possible when the communications infrastructure supporting the applications has the redundancy needed to maintain continuity.

Similarly, it's critical that first responders at state and local governments maintain communications with one another. But they must also retain the ability to send out alerts via social media and to access public records that might come into play during a disaster.

UC also supports COOP when communications extend to remote workers and teleworkers. UC and telework are two sides of the same coin. Telework itself is a COOP strategy. However, it functions only when teleworkers remain connected. One way to ensure availability to teleworkers: Have wireless services on standby for when wired access becomes unavailable.

After intra- and interagency communications and remote workers, the third leg of the UC tripod that supports COOP resides in the cloud. Virtualized communications infrastructure and applications can be more easily replicated by third-party IaaS and SaaS providers. Those providers, as noted in the *Data Access and Storage* section, are hired by a growing number of government agencies as backup and disaster recovery (or even primary) host facilities.

Power Considerations

Data centers consume an estimated 1.3 percent of the electricity used in the United States. It's impossible to separate COOP planning from power management. In the broader definition of COOP as both basic disaster recovery and assurance of high availability and quality of service, power management has two components:

- Delivering backup to and conditioning of electricity from the utility coming into the data center
- Cooling servers and racks within the data center to prevent failures because of overheating

Three tiers of power ensure continuous data center operation: the feed from the local utility, a battery bank for immediate failover in a blackout or storm, and a self-powered uninterruptable power supply (UPS) that takes over from the batteries, typically within minutes or seconds.

The data center staff must optimize each of these power supply tiers. Some approaches to this include IT managers requesting that the utility move power lines from overhead to underground; the utility adding redundant feeds from a different substation; and (where available) the agency signing up with a separate utility.

Connections and switching gear should be tested periodically for reliability and safety. Data center operators also should filter utility power to remove spikes and dips, and they should continuously monitor incoming power to ensure that it's clean.

Similarly, best practices dictate that IT staffs conduct regular tests of the interconnections between the battery and UPS equipment. That includes revving up the generators that power data center–grade UPS systems and regularly changing or adding stabilizers to diesel fuel.

Designing internal cooling systems has become something of an art as traditional mainframes have given way to towers, racks and blade systems. Government mandates for the greening of all federal facilities, with special emphasis on data centers, also plays a role. The challenge for data center operators becomes how to provide sufficient cooling for COOP assurance in a way that maximizes each watt of power.

Modern cooling emphasizes the use of passive convection methods where possible. Lower-power ARM chip servers now hitting the market promise greater use of passive cooling, which uses no electricity, or just enough for air exchange fans.

For dense racks and blade groupings that require more powerful cooling, most IT staffs still want to maximize efficiency by applying cooling directly where it's needed, such as the aisle between rows or racks situated so that the hot ends of equipment are back to back.

Keep in mind that while maximum efficiency and minimum electricity use are important goals for IT managers, equipment failure and loss of application availability because of insufficient power and cooling are unforgivable for critical online services.