CBI for the Cloud

By Lon Berk[1]

Originally Published in the American Bar Association Committee on Insurance Litigation Coverage Newsletter - Vol. 21, No. 6

Business computing is being transformed by the cloud. Rather than maintaining and controlling their data and computation resources on site, with increasing regularity businesses are relying upon third-parties to provide and maintain these resources. Under this model, access to hardware, software and applications are provided to businesses through internet connections. A vendor, such as Amazon.com, permits the client to connect to a server farm providing the computation resources the client needs, as needed. Consequently, businesses can avoid the expense of acquiring and maintaining their own information technology and essentially rent time on the vendors' servers. Through virtualization, a programming technique that effectively permits a single server (or server group) to appear and operate as multiple servers, multiple businesses can operate their own computation systems on a single server system, obtaining the benefits of owning their own hardware but in fact acting as tenants on hardware belonging to another party.

The frequency of cloud computing is expected to increase as its advantages over individual ownership of a business' computation infrastructure are evident. As one commentator describes it:

> The shift toward cloud computing is driven by many factors including ubiquity of access (all you need is a browser), ease of management (no need for user experience improvements as no configuration or back up is needed), and less investment (affordable enterprise solution deployed on a pay-per-use basis for the hardware, with systems software provided by the cloud providers). Furthermore, cloud computing offers many advantages to vendors, such as easily managed infrastructure because the data center has homogeneous hardware and system software. Moreover, they are under the control of a single, knowledgeable entity.[2]

As the incentives are shared by both the cloud computing vendors and users, the trend towards cloud computing should be expected to increase. Businesses save the expense of acquiring hardware; while vendors are able to concentrate their skills on developing appropriate systems to meet their customers needs.

There are risks attendant with this change of computing paradigm. As a consequence, critical operations are conducted off-site, in areas that the business enterprise may not control, and are subject to risks that cannot be evaluated with attention focused merely on onsite business operations. Moreover, businesses operating under this model are subject to risks that cannot be evaluated merely by a study of the geography in which the company operates. Server farms providing computation services may be located in jurisdictions far from the actual location in which the company operates. Much attention has

[1] Lon Berk is a partner at Hunton & Williams LLP and a member of the firm's Insurance Litigation & Counseling group. The views expressed in this article are neither the firm's nor its clients'

[2] Gillam, Lee, *Cloud Computing* at 23-23.

focused on privacy and cyber-security issues arising out of this transformation. The use of third-party vendors for data retention and computation, some have contended, increases the vulnerabilities of hackers acquiring a company's proprietary information. Whenever access over the internet is required to access data, there is a risk that unlicensed third-parties may access it as well. Given the already wide-spread permission by businesses of off-site access to computation infrastructures, it is not clear that the move to cloud computing will increase the overall risk of such data breaches. There is in fact the possibility that having data and computation maintained by those committed to its preservation, security may actually be increased, rather than decreased, by cloud computing.[3] This article focuses not on security issues relating to data breaches, but on a perhaps related risk arising out of cloud computing — business interruption. With companies giving control of the maintenance of their data and computation needs to third-parties, their business continuity depends upon the continuity of that third-party. In exchange for a reduction of the cost of maintaining a computation infra-structure, the business is accepting the risk of depending upon a third-party's operations. It can be expected that companies will look to their insurers to cover this risk. In connection with security breaches, insurers have taken the position that data is not "tangible property" and, in response, have denied coverage under their standard comprehensive general liability forms, with some incorporating exclusions in their standard forms and offering to sell coverage under other forms designed to cover cyber-security events. Although to avoid litigation, it may be prudent to acquire such specialized coverage, insureds should be aware that business interruption losses resulting from loss of access to the cloud should, in the majority of cases, be covered under so-called "legacy" contingent business interruption forms as well.

Business interruption in the cloud

This risk just recently was demonstrated. On April 21, a portion of the Amazon Elastic Computer Cloud ("Amazon's EC2") was disrupted, interfering with the business operations of its customers. Amazon's EC2 permits users to run EC2 instances[4] — essentially, a virtual computer in the cloud — that can be used as a user would use, say, an ordinary laptop. That is, the instance has the operating system, software and applications desired by the user. Generally, when these instances are terminated, their data is eliminated, as the instance vanishes. However, Amazon's EC2 gives customers the option to create an instance from its Elastic Block Store ("EBS"). The EBS permits read and write access from EC2 instances, acting essentially like a virtual "server" for all the EC2 instances the customer uses so that information can be stored and exchanged. Using the EBS, a user can stop an instance, and have its data stored in the EBS. Thus, by using the Amazon EC2 service with the EBS, the customer can dispense with investments in servers, relying upon the Amazon cloud for computation and data-storage needs. Consequently, when on April 21, Amazon's EBS system was disrupted, customers using the service, were unable to access data and to perform computations and suffered an interruption of their business. This was not the only occasion where dependence on the cloud resulted in such interruptions. More recently, in August an electrical storm struck a server farm in Ireland disrupting cloud access to multiple European companies. As businesses increasingly look to the cloud to satisfy their computation and data needs, we can expect more and more instances where businesses are interrupted as a result of

[3] A significant proportion of data breaches result from social engineering techniques, through which outsiders may gain access to passwords and user names from unwitting employees. That risk is still attendant with cloud computing; however, cloud computing vendors, may be more alert to the risks of such techniques.

[4] An instance is essentially a software object of a specified type. Thus, an EC2 instance would be an object with specified data and methods required by the EC2 type.

interruptions in their access to the cloud. It is therefore imperative that companies understand whether the risk of such interruptions are covered under their insurance programs.

Contingent business interruption coverage

The most promising coverage for losses caused by events such as the interruption of Amazon's web service is that provided by contingent business interruption insurance, or "CBI". CBI is intended to provide coverage for the insured's lost profits caused by loss to a key customer's or supplier's property, or to other so-called "dependent property," not owned by the insured, but relied upon by the insured for its business operations. Typically, CBI provides coverage requires the insurer

> to pay for the actual loss of Business Income the insured sustains due to the necessary suspension of its operations during the "period of restoration." The suspension must be caused by the direct physical loss of or damage to "dependent property."

Often, this coverage is only provided where the dependent property is identified in a schedule and where the cause of loss is identified as covered.[5] "Dependent properties" generally include contributing locations, recipient locations, manufacturing locations and leader locations, where a "contributing location," is the location where the supplier operates, a recipient location is the locations where customers operate, a manufacturing location is where products are manufactured and a leader location is a magnet location that attracts customers to the insured's property. An ISO form, for example, defines "dependent property" to mean:

> property operated by others whom you depend on to:
>
> a. Deliver materials or services to you, or to others for your account (Contributing Locations). With respect to Contributing Locations, services does not meant water, communication or power supply services;
>
> b. Accept your products or services (Recipient Locations).
>
> c. Manufacture products for delivery to your customers under contract of sale (Manufacturing Locations); or
>
> d. Attract customers to your business (Leader Locations).

The key concept unifying all is that the location, although necessary to the insured's profitable operations, is neither owned by the insured nor operated by the insured. It seems evident that the server farms owned by cloud vendors constitute either manufacturing or contributing locations. It is at those farms that the insured may perform critical computation activities and it is from those locations that the insured receives its capacity to perform those activities.

[5] This language is adopted from the Insurance Service Office's Business Income from Dependent Properties — Broad Form.

In *Archer-Daniels Midland Co. v. Phoenix Assurance Co. of New York*,[6] for example, the insured sought CBI coverage for losses resulting when the Mississippi River flooded farmland, causing increased costs of raw materials and transportation. The insured, Archer-Daniels, processed farm products and sustained loses caused by damage to farm property owned by farmers that sold grain to Archer-Daniels' suppliers, as well as damage to the Army Corps of Engineers, which provided navigable waterways. The court agreed. Even though there was no direct contractual relationship between the insured and the Army Corps of Engineers or the farmers, the Court held that because these persons provided services and products upon which Archer-Daniels relied, damage to their property was covered by CBI.

By the same reasoning, server farms used in cloud computing should constitute manufacturing or service property covered by CBI. Damage to that property can directly lead to an interruption of the insured's business by preventing the insured from performing critical computation operations.[7]

Perhaps the most difficult aspect of triggering CBI coverage is the requirement that the dependent location sustain "physical loss or damage." Insurers have argued and some courts have relied mistakenly on those arguments to believe that damage to software or data does not constitute "direct physical loss" and therefore interruptions of server farms resulting from software issues and leading to lack of access to data, arguably, should not be covered. Data or software, it is contended, is not tangible or physical and therefore cannot cause direct physical loss. Similarly, it has been argued, because data is intangible and cannot be touched, damage to data cannot constitute "direct physical loss." Part of the motivation underlying these arguments is the fact that the physical hardware of a server might not be harmed even though data stored on it is no longer retrievable and software used on it is no longer operable and that, if this data and/or software were replaced, say, the server would operate as before. Although there is some initial appeal to such a position, the contention that there is a distinction between physical damage and damage to data and/or software, when looked at more closely, derives from a misunderstanding of the computer operations and is an error.[8]

In fact, insurers have long recognized that software damage is physical damage covered by their first-party property. At the end of the 20th century, there was concern about whether the apparent widespread use in software of two-digit fields for years would cause business interruptions and other calamities. Insurers, concerned about their exposure under first-party policies for such losses incorporated an exclusion on their policies precluding coverage for "loss or damage caused directly or indirectly by "[c]omputer application software" or "[c]omputer operating systems and related software": "due to the inability to correctly recognize, process, distinguish, interpret or accept one or more dates or times."[9] The

[6] 936 F. Supp. 534 (S.D. Ill. 1996)

[7] One might contend that the customer "uses the cloud providers servers and that therefore these servers should be treated not as "dependent property" but as the insured's property, and cite *Zurich Am. Ins. Co. v. ABM Indust., Inc.,* 397 F.3d 158 (2d Cir. 2005). Such an argument appears to stretch the reasoning of the *ABM* decision and does not seem consistent with the cloud computing model. Under these models, the customer gives up the ability to control and own the servers, relying on property of another to perform its computation and data-storage needs.

[8] It is an error easily made. In fact, this author initially took such a position. "Not Always Obvious," 11 Coverage 19 (September/October 2001).

[9] Insurance Services Office, Inc. Form IL 09 35 08 98.

exclusion, at least implicitly, recognized that there was some risk courts would find — and policyholders would expect — policies to treat losses resulting from software or data issues to constitute as "direct physical loss or damage" triggering obligations under these policies.

In fact, in other contexts, courts have generally found that software and data is tangible, material property. For example, the United States Copyright Act defines "copies" as "material objects … and from which the work can be perceived, reproduced or otherwise communicated either directly or with the aid of a machine or device."[10] The issue has arisen as to whether by loading computer software from read only memory to random access memory the copyright infringement occurs. The general rule is that it does. Because random access memory is temporary and eliminated once power to the computer is turned off, defendants have argued that the copy is not a "material object[]" "from which the work can be perceived…." But this argument has been rejected. For example the Ninth Circuit has held:

> [I]t is generally accepted that the loading of software into a computer constitutes the creation of a copy under the Copyright Act.…[S]ince we find that the copy created in the RAM can be "perceived, reproduced or otherwise communicated," we hold that the loading of software into the RAM creates a copy under the Copyright Act.[11]

Nonetheless, in the insurance context, courts have found otherwise. In *America O-Line, Inc. v. St. Paul Mercury Ins. Co.*, [12] for example, coverage was sought for a class action alleging that AOL's 5.0 internet access software disrupted internet and local area network connections, caused computers to crash and corrupted systems and files. The court addressed whether computer data and software were tangible property. It concluded that they were not reasoning

> Computer data can be transmitted and sorted in a variety of ways, but none of them renders the data capable of being touched. A "bit" on a computer disk or hard drive is not palpable. Electrical impulses that carry computer data may be observable with the aid of a computer, but they are invisible to the human eye. [13]

The court concluded that data and software were therefore not tangible. Although the Court did not expressly address whether damage to data or damage caused by software constituted direct physical loss or damage, the decisions reasoning can be used to support the view that it is not and that CBI coverage therefore does not apply.

This conclusion was reached in *Ward General Ins. Services, Inc. v. Employers Fire Ins. Co.*[14] There, an error while the insured was attempting to update its data base corrupted data, forcing the insured to re-

[10] 17 U.S.C. § 101.

[11] *MAI Sys. Corp. v. Peak Computer, Inc.*, 991 F.2d 511, 519 (9th Cir. 1993).

[12] 207 F. Supp. 2d 459. (E.D.Va. 2002). *America Online* concerned a general liability policy, but its reasoning would nonetheless be applicable to the issue under CBI.

[13] 207 F.Supp.2d at 467.

[14] 7 Cal. Rptr. 3d 844 (Cal. Ct. App. 2004).

input data. The insured sought to recover the costs of recreating its data. The court, however, found that loss of data does not constitute a physical loss unless the medium on which the data is stored is itself injured. Because the insured was seeking to recover the expense of re-inputting its data, not the replacement of the hardware on which the data was input, the court found that there was no physical loss.

As mentioned, there may be something initially plausible about the reasoning of these decisions — after all, it matters not whether our data is stored on this hardware or that, so long as it can be accessed when we want it. But the reasoning is fundamentally flawed and arguably should not be followed, especially where CBI is at issue.

To see the source of the courts' errors, we need to consider on a general level what a computer is and does. Essentially a computer can be seen as just a device that has exceptionally many on/off switches and manipulates them exceptionally quickly. When data is lost, it is because the state of these on/off switches has been lost. That is a fundamentally physical loss, not intangible one, even if the way we might repair that loss is to reset the switches into the configuration they were earlier in. When software fails it is because the manipulation of these on/off switches does not occur as planned. That is fundamentally a physical loss even though we might repair it by modifying programming instructions, which ultimately are themselves just configurations of the on/off switches.

Imagine a device with eight slots that can contain one marble each. If a marble is in a slot, we say the slot is on; and, if the slot is empty, we say it is off. For ease of reference, let's call the first slot "0", the second "1", and so on, with the eighth slot called "7". We can physically design this device, perhaps using spring releases so that if a marble of a certain weight is dropped into a slot that contains a marble, both marbles will pop out of that slot with one of them dropping into the next highest numbered slot, if such a slot exists. So, for example, if we drop a marble into slot 5, the marble will stay there if slot 5 was empty, and, if there was already a marble in slot 5, the new marble will pop into slot 6 and the old marble discarded from the device. Then, if slot 6 was empty, the marble will stay there, and otherwise, the new marble in slot 6 (that is, the old one from slot 5) will be dropped into slot 7. If slot 7 has a marble, both will be dumped from the device. Let's call this device a "marble byte."

We can use marble bytes to do simple arithmetic and multiplication as well as other arithmetic operations provided that the result of these operations is sufficiently small (in particular not greater than 255). For instance, say we have three marble bytes, call them byte 0, byte 1 and byte 2. Then, if byte 0 has a marble in its slot 1, with all other slots empty and byte 1 has a marble in its slot 1 with all other slots empty and we take those marbles and put them (one at a time) into byte 2's slot 1, the result will be that byte 2 will contain a marble in its slot 2, while all its other slots are empty. We can take this result to be a computation of 2+2=4. Similarly, if we put a marble in slot 0 and slot 3 of a marble byte, and simply leave the device alone, we can take ourselves to have stored the number 9.

We might, moreover, combine more marble bytes together in a fashion that has marbles leaving the slot 7 of one device to go into the slot 0 of another. For instance, we might put 16 of these devices together, each device labeled "0", "1", "2", "3", "4", "5", "6", "7", "8", "9", "A", "B", "C", "D", "E", and "F" and set them up so that when a marble drops out of slot 7 of device 8, that marble is dropped into slot 0 of device 9, and similarly, when a marble is dropped out of slot 7 of device 9, a marble is dropped into slot 0 of device A, and so on. We can call these conglomeration of marble bytes, "marble words" and store significantly higher numbers using marble words and perform more complicated arithmetical operations.

Now, if something were to go wrong with our marble word so that it failed to pass marbles as expected, we would conclude the device has been physically damaged. In other words, if we found that marbles did

not go from marble byte A's slot 7 to marble byte B's slot 0, we would conclude there was a physical problem with the device, and we would have to do something physically to it to fix it so that the marbles were transferred correctly. And if somehow or another, we had stored the number 9 in a marble byte as described two paragraphs back and then could not find a marble in any of the slots, we would conclude that the device had sustained a physical loss; that it had lost its marbles. If we desired to fix it, we would replace the marbles as before – that would be physically repairing the marble byte.

We could eventually put as many of these devices together as we wanted and construct different sorts of mechanisms for having marbles transfer from slots to slots and bytes to bytes. If we were, for instance, to put together one billion of these marble bytes, we would have created a giga-marble byte, in effect a hard-drive with one gigabyte of capacity, for treating our marble slots are each a bit, with eight of them forming a byte and the 16 forming a word. We can also imagine that if certain slots contain a marble, the devices method of transferring marbles is changed. For instance, say we have marble bytes 0, 1, 2 and 3 and that the device is designed so that if byte 3 contains no marbles, bytes 0,1 and 2 operate as described before; but, if a marble appears in slot 1 of byte 3, then when a marble is put into slot 7 of byte 1 and that slot already contains a marble, then one marble is transferred to slot 0 of byte 0 and another is transferred to slot 0 of byte 2. Different configurations of bytes, would, in this device, lead to different methods of transferring marbles from slots to slots. In principle, we could create an enormous marble device that is programmable and that could do (although much much much slower and using much much more space) any computation that could be done by a modern digital computer.[15]

Now if something happens to the marble device so that it fails to shift marbles to the correct slot, or if it drops a marble from a slot when it shouldn't, as we mentioned, there would be little doubt that the device sustained physical loss or damage and we would repair the device by taking physical steps, such as putting marbles in the correct slots. By the same token, if something should happen to a computer so that it fails to shift charges or drops a charge when it shouldn't, there should be no doubt that the computer has sustained physical loss or damage as well. The fact is that when marbles are put into the correct slots of the marble device, we are making a physical repair, and, by the same token, when an electromagnetic charge is put in the right spot of a computer, we are making a physical repair. This is true even though in one case we may be dropping marbles in a slot while in the other we are typing code into a console.

What may complicate the analysis of this issue somewhat is the fact that, unlike the marble device, we are able to load instructions onto a computer that affect the shifting of charges from one portion of the disk to another. We are able in other words to program a computer so that it might perform (subject to limits on memory and time) any computation possible. It therefore appears that, when a computer fails to perform a shift of charges correctly, it is not because of a physical loss or damage, but because of an intangible failure to comply with instructions, but that appearance is just an appearance. Ultimately, the failure is a failure not of instructions but of the shifting of charges, that is, of an actual physical event. When code is input into a computer, it is "compiled" into machine language into a sequence of 1's and 0's of high and low electric charges. The computer that fails to perform due to a software failure, therefore fails to perform physical operations, just as the marble device that fails to shift marbles, fails to perform a physical operation.

[15] This is a consequence of Church's thesis, which roughly is that any type of computation device developed will perform precisely the same set of computations as any other type.

It is the failure to appreciate that there is little principled difference (ignoring size and speed) between what we can do with marbles and what we can do with charges that may lead to decisions such as *America Online* and *Ward General* which attempt to distinguish between physical loss and data loss. Other courts (and commentators) have not made the same mistake. *American Guarantee & Liability Ins. Co. v. Ingram Micro, Inc.* [16] is one example. The issue before the trial court was whether a power outage caused "direct physical loss or damage" under an all risks legacy first-party policy. The Court described the against coverage as follows:

> American and its expert witnesses admit that Ingram's mainframe computers and the matrix switch did not function as before the power outage and that certain data entry and reconfiguration processes were necessary to make Impulse operate as it had before the power outage. American argues however, that the computer system and the matrix switch were not "physically damaged" because their capability to perform their intended functions remained intact. The power outage did not adversely affect the equipment's inherent ability to accept and process data and configuration settings when they were subsequently reentered into the computer system.[17]

It rejected this argument, reasoning that:

> At a time when computer technology dominates our professional as well as our personal lives, the Court must side with Ingram's broader definition of "physical damage." The Court finds that "physical damage" is not restricted to the physical destruction or harm of computer circuitry but includes loss of access, loss of use, and loss of functionality.[18]

Just as the loss of the marble device's functionality — its inability to properly shift marbles to a higher numbered slot — is "physical damage", so too should be a computer's loss of functionality — its inability to properly shift electromagnetic charges.

Similarly, in *Southeast Mental Health Center, Inc. v Pacific Ins. Co., Ltd.*,[19] the court rejected an insurer's contention that there was no coverage for its insured's loss of business income resulting from the corruption of data on a computer. Hurricane Elvis destroyed power and utility poles, causing the insured to lose electricity and telephone service, which, according to the insured, damaged its computer causing the loss of data. Consequently, the insured lost significant business income. The insurer contended that the loss of data did not constitute a physical loss covered by the policy.[20] The district court disagreed, holding that "corruption of the [insured's] computer constitutes 'direct physical loss of or damage to

[16] 2000 WL 726789 (D.Ariz. April 18, 2000).

[17] *Id.* at 2.

[18] *Id.*

[19] 439 F.Supp.2d 831 (W.D. Tenn. 2006).

[20] *Id.* at 837.

property' under the business interruption policy."[21] The court found, citing the *Ingram Micro* reasoning, that "The computer[] 'physically lost programming information and custom configurations necessary for [it] to function' when [it] was damaged by the power outage," entitling the insured to business interruption coverage.[22]

This result is supported by the reasoning of *Wakefern Food Corporation v. Liberty Mut. Fire Ins. Co.*[23] There the insured operated supermarkets in the Northeastern United States. Its policy insured against consequential loss or damage resulting from an interruption of electrical power caused by "physical damage." In August 2003, there was a massive power outage in the Northeast with "[a]t least 265 power plants with more than 508 individual generating units … shut down," in some parts for as long as four days. As a result, the insured's food was spoiled and it sustained consequential losses.[24] The insurer denied coverage, however, contending that the power lines were not physically damaged. The insurer essentially argued that the blackout was caused by the operation of safety devices that prevented the electrical grid from function until certain safety steps were taken. The New Jersey court held that nonetheless "from the perspective of the millions of customers deprived of electric power for several days, the system certainly suffered physical damage, because it was incapable of providing electricity."[25] The court cited with approval *Southeast Mental Health Center*'s finding that "'physical damage' could include loss of 'functionality' even if the affected machinery remained intact."[26] By the same token, from the perspective of its clients, when a cloud provider is unable to provide necessary data and computation services, the system has suffered physical damage. The central idea is that whether there is physical damage is determined by reference to the viewpoint of the insured. An electric grid that does not provide electricity is physically damaged from the perspective of the insured. The grid is physically incapable of providing the insured with electricity, even if, after certain safety devices are re-set, the grid will function without further repair. Similarly, a server farm that fails to provide its clients with computation services is physically damaged from the perspective of the insured customer, even if those computation services might be provided by correcting software on the system.

Electronic data limitations

More recently, perhaps recognizing the force of the reasoning that loss of data constitutes a physical loss, insurers have incorporated policy exclusions purporting to restrict the coverage available for data loss and have offered separate policies to cover losses due to data and software corruption. The extent to which these exclusions restrict coverage, however, is a matter that needs to be addressed. The ISO CBI form currently contains the following language:

[21] *Id.* at 837-38.

[22] *Id.* at 838.

[23] 968 A.2d 724 (N.J., Super. Ct, App. Div. 2009)

[24] *Id.* at 728 – 29.

[25] *Id.* at 736.

[26] *Id.*

> We will pay for the actual loss of Business Income you sustain due to the necessary "suspension" of your "operations" during the "period of restoration". The "suspension" must be caused by direct physical loss of or damage to "dependent property" at a premises described in the Schedule caused by or resulting from a Covered Cause of Loss. However, coverage under this endorsement does not apply when the only loss to "dependent property" is loss or damage to electronic data, including destruction or corruption of electronic data. If the "dependent property" sustains loss or damage to electronic data and other property, coverage under this endorsement will not continue once the other property is repaired, rebuilt or replaced. The term electronic data has the meaning set forth in the Coverage Form to which this endorsement applies.[27]

"Electronic data" is often defined as:

> information, facts or computer programs stored as or on, created or used on, or transmitted to or from computer software (including systems and application software), on hard or floppy disks, CD-ROMs, tapes, drives, cells, data processing devices or any other repositories of computer software which are used with electronically controlled equipment. The term computer programs, referred to in the foregoing description of electronic data, means a set of related electronic instructions which direct the operations and functions of a computer or device connected to it, which enable the computer or device to receive, process, store, retrieve or send data.[28]

This exclusion should not apply to the vast majority of incidents that might result in interruption of computation services provided by cloud vendors. First, note that it is drafted in terms of the "dependent property" sustaining a loss or damage to electronic data. In the vast majority of cases, it will not be the dependent property that sustains such a loss, but the insured's property that is unable to access data at a vendors server farms, that is, at the dependent property. The data in other words may remain intact at the server property, but not accessible by the customer because of other loss at the dependent property. In the recent example of an electronic storm damaging a server farm so that European customers were not able to access it, this exclusion plainly would not apply.

As another example, consider the April 21 Amazon Web Services incident described several pages back.[29] The portion disrupted was the EBS. Recall that the EBS permits EC2 instances to read and write data so that it can be used after instances are stopped. The system works by creating a cluster of EBS nodes to which the ECS instances store data and serve requests. EBS data is replicated to multiple nodes to increase availability and durability. The nodes are connected to each other and replicate data. But when

[27] *See* ISO Form CP 15 08 04 02..

[28] *See* ISO Form CP 00 10 04 02.

[29] The following summary of the causes of the interruption is based upon a report from Amazon Web Services published on its website.

a node loses connection to one to which it is replicating data, it assumes the other node failed and searches for a new node to replicate its data. In order to upgrade their system, Amazon redirected certain requests from the instances. Unfortunately, the traffic shift was performed incorrectly and, as a result, nodes in a portion of the EBS cluster were isolated from other nodes in that cluster. Nodes then began to search the cluster for additional server space for their data. But the nodes got stuck in a loop. The result was a "re-mirroring" storm that spread to other parts of the system, forcing a significant portion of the system to become "stuck." Among the steps taken to respond to this "outage" was adding additional storage capacity so that the stuck volumes could replicate their data. Thus, the response to the event was a physical response; that is to add new capacity to the system. That response, coupled with the discussion of the last section, strongly suggests that the AWS outage was the result of a loss of physical storage capacity, which is loss, not of an intangible, but of tangible physical space.

In sum, CBI forms may be triggered by many, if not all, interruptions of access to the cloud and are a source of coverage for resulting business interruption losses. When confronted with such losses, insured's needing coverage should examine the nature of the loss and look to their "legacy" CBI policies, as well as any new specialized policies they may have.