DISASTER RECOVERY / BUSINESS CONTINUITY

At this 01.31.23 Roundtable 55, *NOREX Members discussed practicing / simulating DR and on what schedule; process to determine critical systems for restoration; DR strategy moving from on-prem to Cloud; separating backups from production to protect against ransomware events; steps to review suitability of DR plans for a ransomware attack; mass communications systems in use to notify stakeholders of major incidents; prioritizing processes / applications; and ultimate accountability for DR / BC.*

EXECUTIVE SUMMARY

A Member asked, when practicing Disaster Recovery do you simulate a DR and on what kind of schedule? An Enterprise Security Officer shared that they do a quarterly DR test, swinging all necessary systems for production over to the DR site for the weekend. They run user acceptance testing then swing things back. They also pick one system of increasing complexity to do a complete rebuild from backup, assuming a worst-case DR situation. They have been doing quarterlies for many years, and it's proved quite useful. A Head of IT Operational Control shared they do a full data center failover. All their applications and infrastructure are part of this. Three years ago, they had a DR consultant from their external auditors do a full data center failover as part of this process and then fail back from that test. A Director, Global Continuity, Resilience, and IT DR shared they do failover between regions from UK to US, UK to Canada, and vice versa. They do a full failover and a partial failover, but he feels they do too many tests a year. They have 3,000 applications to test a year besides the infrastructure. He feels they have to improve and automate as much as possible.

On the topic of how DR strategy and approach changes as organizations move from on-prem toward more Cloud-based, an Enterprise Security Officer on the East Coast stated that when Hurricane Sandy hit, they migrated their DR environment into the Cloud. They did not want to be dependent on an East Coast power grid with their international clientele. With a physical data center, it was both expensive and non-feasible to have a model of rebuilding an application from off-site air gapped backup as a part of their general process. Now it is a straightforward thing they can aspire to improve. There are pricing issues, but the Cloud has made disaster recovery much better. A Director of Global Continuity, Resilience, and IT DR stated that they are attempting to be 60% Cloud in the next five years. He finds that it is near impossible to achieve this switch to Cloud because of the huge number of apps and regions they have. They are expanding their data centers and he believes it will not go smoothly because of the cost, the way it is managed in the Cloud, and the regulations they have on every continent. A Senior Manager Infrastructure Services shared they first had to unravel DR / BC with the business side. There is always a conflict of interpretation. They found difficulty in marrying up their on-prem requirements with their ability to scale to the Cloud. The technology was ahead of the actual business technology. There were many on-prem devices that needed a connection or had GPOs that keep them alive and running. The strategy for them is to slow down. It was initially a three-year migration plan that moved to five to six years, because there are things that cannot be unraveled. It is just something you don't know until you don't know.

Additional headline topics:

- **Mass communication systems** in use for incident / disaster notification.
- **Prioritizing** processes and applications.
- **Determining which systems** are critical for restoration.

TABLE OF CONTENTS

NOREX Roundtable 55 Transcript
Disaster Recovery / Business Continuity
January 31, 2023

Moderator: Good morning, everybody. I'll be your moderator this morning for our conversation around disaster recovery and business continuity. The first question today is around practicing your disaster recovery. Do you simulate the DR? Mary, if you wouldn't mind kicking that off. I've just put up a poll. As far as scheduling, do you do it once a year, twice a year, whatever? Take a minute and respond to the poll. Mary?

TOPIC: Schedule of DR plan testing

POLL: How often do you test/practice your DR plan:

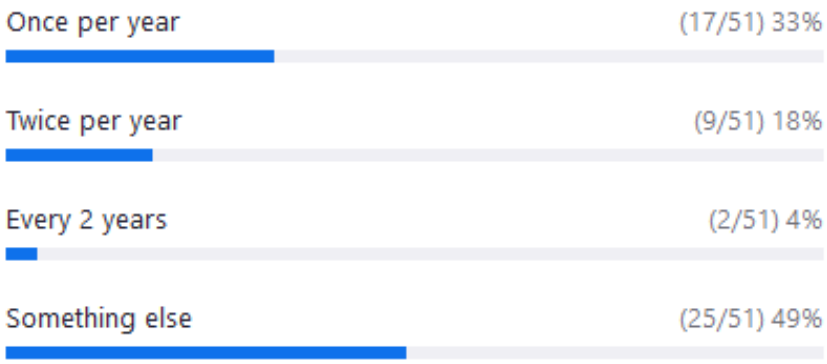

POLL: Are your DR test plans full or partial:

Mary W.: Good morning, everyone. I don't know how much more detail I can add to it. But just curious what other folks are doing. We do have a schedule that we kind of try to adhere to, practicing our disaster recovery for our multiple applications. We do them application by application, not a wholesale, everything's down practice. Just curious what everyone else does, thanks.

Stephen T.: We do a quarterly DR test, swinging all necessary systems for production over to the DR site for weekend and run user acceptance testing then swing things back. We also pick one system of increasing ugliness and complexity to do a complete rebuild from backup, assuming a worst case DR situation ran somewhere something where you have to rebuild from an old scratch backup on a completely new virtual world. That's been our process. We've been doing the quarterlies for many,

many years quite happily and it's proved quite useful. The business is getting progressively more and more involved in the process as we actually get positive leadership who are invested in maintaining business resilience.

Matthew W.: DR test quarterly, backup & DR replica testing weekly (Veeam SureBackup).

Brandon S.: How big are your IT departments?

John K.: 10.

Bryan K.: 9.

Joshua K.: 94

Jim J.: Two IT, five devs.

Scott D.: Four desktop, three network, six devs.

Niki S.: 37.

Stephen T.: 70. We fail from NYC to AWS west coast.

Robert M.: Three desktop, two network / infrastructure (myself and director), four apps / dev, one telecom.

Tom Q.: Replying to "How Big are your IT..." We have a distributed model with about 300 endpoint / network / data center / development.

Natasha R.: For our disaster recovery we do a full data center failover. That's all of our applications, all of our infrastructure as part of that. About three years ago, we did a business unit by business unit level. We had a DR consultant from our external auditors, which actually recommended doing a full data center failover as part of that process. We switched up over to a full data center failover itself over a weekend and then fail back from that test. I know that from obviously some of the regulations coming in just from operational resilience from the UK that they will start asking scenario-based testing on that. We're just working through that as well as doing the full failover. But that's from sort of the UK side.

Moderator: Thanks, Natasha. Jeff, you have a comment?

Jeff M.: I have a couple of questions. Number one, when you're talking about full failovers, how far apart are your data centers? How close are the data centers in proximity for full failover?

Matthew W.: Mine's in Texas and we are in Illinois.

Stephen T.: Quarterly failover of all production applications. Rebuild from scratch for key applications one at a time each testing cycle.

Jim J.: We're in Kansas, DR site is in Houston, TX.

Robert G.: We're in Salt Lake and data center is in the Boston area and in northern VA.

Matthew W.: We have a matrix of what fails what. Our cyber insurance requires that including how fast we can get back if we get ransomware, etc. RTO / RPO.

Stephen T.: I would suggest Infrastructure as data is a key step.

Paula B.: If you are able to share any documents with NOREX, please put a note in the chat and we will follow up with you. Thank you!

John K.: We cut off the primary data center.

Matthew W.: So do we. We either split it with Veeam Network Appliance rerouting for one-to-one VM failover or have everyone just VPN into DR site.

Gregory R.: Ours is standalone as well. We're figuring out the Active Directory situation at the moment. AWS AD wasn't cutting it. Forgot to mention the "DR" is in AWS.

John K.: First, do they have PII? We are an Okta shop too. Failover with Okta agent was seamless the last DR test.

Matthew W.: Same. We have the agents on DCs at the remote sites too.

Lauri H.: Bryan, how did you get your list for the cards?

Bryan K.: From our initial inventory of systems and services. We spent a few months winnowing the list down from 80 to 50 so that it would be more manageable from an exercise. Our service management system was a useful dataset for type of services / systems.

Natasha R.: We have completely different locations. One's in central London, and one is outside of central London for that. Completely separate data centers and different actual vendors that own the different data centers itself from testing the different sites completely.

Jeff M.: Okay. The reason why I ask that is because I've been in companies where the data centers were less than 100 miles apart, so full failovers really weren't impactful at all. But with the current company that I'm in, our data center is about 1100 miles from each other. Needless to say, latency has a big impact on being able to fail those over. From our perspective we do failover application by application. We have a select group of applications which we consider financially significant, and those are done approximately once a month. We pick between 10 and 12 per year to do. The challenge with that is a lot of them that we pick to do there are so many integrations trying to establish which ones we pull across and which ones we don't, because if we don't pull them all across to the other data center then the latency impacts the testing.

Gregory R.: We have an interesting DR. It's a sliver of the production environment. I think Jeff kind of hit it, the most economically important part of the business is DR. It was unique because it was very finite amount. They picked one workflow out to develop for DR and the rest sort of got ignored. It was interesting way to go about it. I think now there was success, and it wants to be expanded. There's a very verbose written policy, but no data governance or anything. There's pieces that are complete and other parts there. I just wanted to share a partial scenario.

Ammar Y.: We do failover between regions as well from UK to US, UK to Canada, and vice versa. And of course we do the failover in the UK as well. In each region we do our testing and believe we do it between regions. Now, regarding latency, yes, it's an issue, but we do have to do it. For the

scenario-based, this is what we are introducing. We do a full failover. We do like partial failover, but we do too many tests a year. This is a problem that we are trying to solve. We have like, maybe about 3,000 applications to test a year, so it's a huge number besides the infrastructure. I'm talking just application to application. That's what we do. And yes, do we do it the best method? Not yet. We have to improve and have to automate as much as possible.

Gaurav J.: Just a question for the group here that we are talking about DR and the failovers. But does anyone in their organization have a documented procedure for the different failovers and how frequently are those testing happening? Do you guys work based off on a plan, like are those procedures and processes in place in a documented format?

Ammar Y.: I can jump in. I've been doing this for over 22 years, the same job, DR, business continuity. Yes, we do have procedures, we do have plans. We are right now working on even creating like a playbook for data centers, small data centers and big data centers. A viable one, of course. We have to be careful that there must be some changes. I think one of the participants said every 3 months, every quarter we check that we do have the maintenance program for that as well to make sure that plans are up to date. There will be a hit and miss, of course, but there is a maintenance program for that. The procedures are there, plans are there for almost everything. Now we are trying to refine it even more. Infrastructure is always kind of a headache for us, because this is a huge environment.

Gaurav J.: Thanks Ammar.

Natasha R.: And I'd say from my side we have the same thing. We have plans in place. I think the thing that we struggle with is the constant change of our infrastructure, obviously new systems being added, technology is changing, etc. I'd be interested to hear how other people handle their DR plans when you have a constantly changing infrastructure, and how you keep your plans up to date.

Mackenzie F.: More of a question for those with DR plans that they've been testing for quite some time. We actually just came out of a scenario where we had to enact our DR plan, and a lot of good things came out of that, and a lot of lessons learned came out of that as well. But my biggest question is, and something that's come up a few times. When you guys do your DR scenarios, is it simply moving those application servers, whatever the case may be, to the other side? Or are you going as far as to actually completely cut off the data center to test a real world if it were to go down?

Stephen T.: We actually block our direct data center access to stimulate a full data center down without actually powering things off. I commend to people the notion of infrastructure as data. If you have a dynamic world, it's the only way to manage it sensibly. And DR is just another aspect of what you get out of the management.

Christian P.: Regarding the previous question about when you have an ever-changing environment, how do you ensure that new systems are conforming to what you need for DR? Much of this needs to be done in the pre-planning exercises and activities that you have set up for any sort of onboarding of a new system or application. One of the strategies that we use is we've got a well-developed tier ranking system that shows the criticality of a system when it comes to corporate survivability or customer impact. We start off with a Tier 0 level, which would be base level infrastructure, without which there is no identity. There is no way of running the business, so Active Directory DMS. We build upon that with Tier 1, which would be those systems that must come up the quickest if we have a disaster. Tier 2 would be more of the deferred systems in terms of maybe if the Tier 1 was 24 hours, Tier 2 would be 48 hours. Then Tier 3, etc., those are deferred as far as prioritization, to get the business back together. If you have that established and you're bringing in new systems and you

can identify with the owners where it meets in that strata, that will help very much when it comes to actually putting things in place.

Bryan K.: Yes to Christian's criticality point! Prioritize in tiers.

TOPIC: Process used to determine systems critical for restoration

Moderator: Okay, very good. Great conversation. I'm going to move on to the next topic. What process is being used to determine which systems are critical for restoration? Anybody want to mention anything on this? While you're thinking about that, I'm going to put up a poll. What process do you use to determine what systems are critical? Do you use a business impact analysis, or is that just formally part of your DR / BC planning?

POLL: What process is used to determine which systems are critical for restoration:

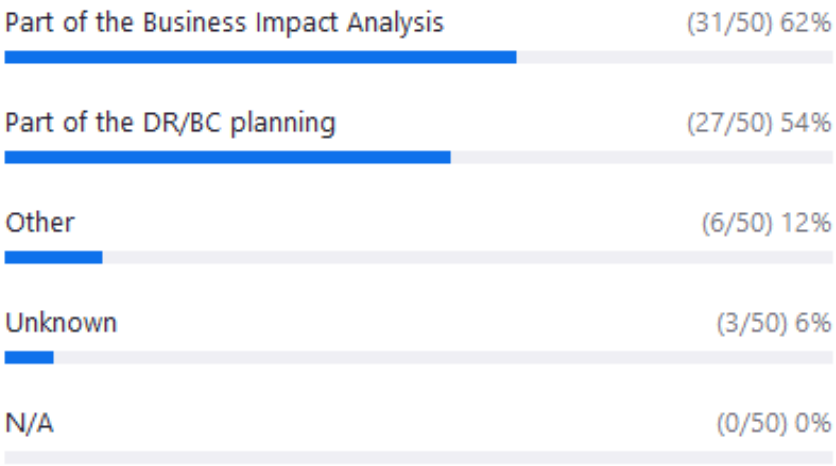

Mackenzie F.: Yes, as I mentioned before, our DR plan was kind of new for us, and most recently enacted. Ideally for us, we have our critical as an IT team. We have our critical infrastructure items that we were able to identify and make a part of that DR plan, and admittedly, even we miss some in that regard. But we go back to our stakeholders, the leaders of various departments yearly, and we present them the previous DR plan or the previous year's. We ask them to identify and update any items that they think may be needed in the future. Admittedly, that in and of itself has some shortcomings there. I'm very interested to hear other people as well.

Matthew W.: Yeah, basically the same thing. We get the stakeholders involved. When we test our DR we make them log into the environment and make sure is everything you expect to see if we're completely down and partially down online in available files, databases, applications, web servers? We're an Okta shop, so we have to get the active agents back up as soon as possible. Those are Active Directory, the VPN, and all of our identity infrastructure. We try to get that back up as soon as possible. That's basically priority.

Gregory R.: The process that we've been using is BRM (Business Resource Managers). It's real early too. Just putting together a rough picture of our application catalog. We really don't know all the applications. We're kind of building it from the ground up. It's really early, but that's who we're using. Actual folks go through and build out the systems.

Bryan K.: We actually just started a process to do an inventory across the entire organization to talk about the criticality of services. We're starting with a card sorting exercise for leaders around the organization to sort what they think that those things are. The most useful aspect that we've gotten out of is just determining where there's a lot of unanimity around certain things, and there's a lot of disagreement about other things That's a given us a prioritized list to say these are the things we need to have some tough conversations about to determine criticality. And these are the things we all agree are low priority and can get further down the list.

Ammar Y.: Our Business Continuity Management Central Team does the RTOs. When we test we do the RTCs. Then we have to create the match between these two. They have their own questionnaires. Of course we do the test, and we do the post-exercise reporting, and then we create the list to match these two. If let's say the RTC is higher than RTO, then we have to go back to the business. But we have to communicate it in between these two departments. Even under my responsibilities, I do have as well the third-party risk management for a DR perspective as well. We do our due diligence on this to make sure that they have their own plans and they have them on RTOs, and when we do the test with them we have to report the RTC as well from a DR and the RTO from a business continuity perspective under the BIA, of course.

Stephen T.: Each of the business units, as part of their planning process and with the assistance of outside consultants, have developed an entire set of key transactions that they have to complete during the course of a DR test. That's the level that gets surfaced to the executive leadership for review of do we have sufficient coverage? During our last major disaster, which was Super Storm Sandy, we got a very good view, and it appears that things were done well. But barring another massive disaster, it's all prayer and planning, and you hope you get it right.

Moderator: Thank you. Anyone else have a comment? Mary or Jennifer, do you have any?

Jennifer T.: I just want to say thanks for all of the comments, and if anyone feels like their I want to say buckets or categories are built out in a more mature way, if that is something that you're willing to share with NOREX, we would definitely be interested in looking at that. Like how you bucket the items. That would be helpful. Thank you so much.

TOPIC: DR strategy changes moving from on-prem to Cloud-based

Moderator: Yeah, good, comment Jennifer. If any of you have that information, please share that with us and we'll make sure that Jennifer gets it. I'm going to move on to the next question, which was posed by Mary and Mike. How does DR strategy and approach change as organizations move from on-prem toward more Cloud-based? I'm going to put up a final poll on where people have their DR sites. Do you have it strictly on-prem or Cloud, or is it a combination or something else?

POLL: Where is your DR / BC site:

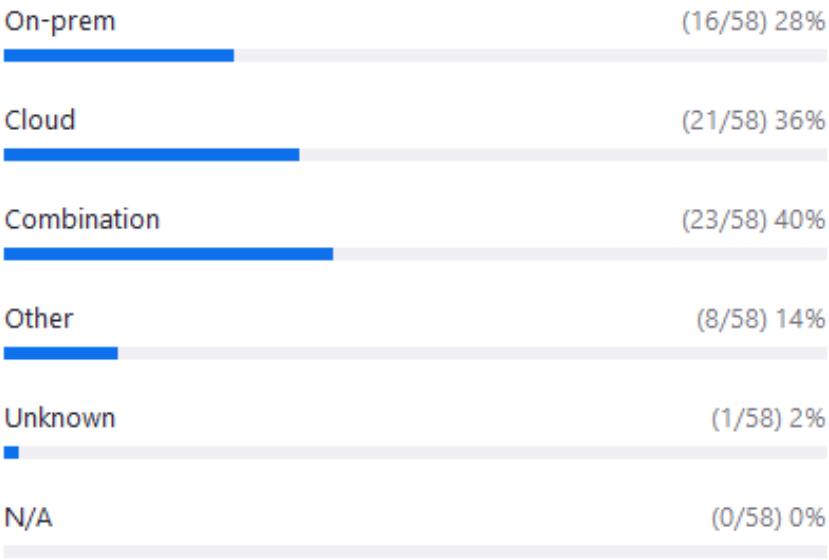

Daniel M.: I feel very out-of-sorts on this conversation. My entire operation is Cloud-based. I have no servers, I have no Active Directory. My process is more business continuity that's determined by my Business Process Steering Committee and the executive team rather than a disaster recovery. If I lose a site, from a technical perspective people just work remotely. I'm very fortunate, but I feel very out-of-sorts, because I don't have a managed DR site. Now I draft my suppliers or my vendors to make sure they're doing their DR, and I get their SOC 2 reports on an annual basis to ensure that they're properly managing their systems. But I've moved everything to the Cloud. I feel out-of-sorts for this call.

Bryan K.: We're in a similar situation, Daniel. We've moved entirely to the Cloud and rely on others' DR for business continuity.

Gregory R.: I need to get me some of that "Business Process Steering Committee". That's awesome, Daniel.

Matthew W.: We use DRaaS (DR as a Service).

Balarabe A.: We are building our DR in the Cloud.

Gaurav J.: We have our DR in MSP's Cloud environment.

Allen S.: We can't move 100% to the Cloud owing to hardware-based security devices, for which there are no virtual equivalents.

Tyler B.: For those that rely on other DR systems, have those companies had to use them? If so, how did that work for your company?

Daniel M.: Allen, what do you mean when you refer to hardware-based security systems?

Matthew W.: For Veeam DRaaS and Cloud Connect, iland, which is now 11:11 Systems which is iland / Sungard Availability Services combined as one company. https://1111systems.com/

Stephen T.: Cloud takes time, but we are in the 5th year of a 3-year migration to the Cloud and should finish it.

Joshua K.: Is anyone using Zerto for their DR solution?

Jim J.: We are using Zerto.

Tyler B.: We use Zerto for BC and DR.

Joshua K.: On-prem, Cloud or both for your targets?

Daniel M.: I have seven staff members, 1000 employees, 40 sites across the US.

Matthew W.: Yay tape robot libraries.

Robert G.: We use Zerto and we like it a lot.

Daniel M.: BC is a business responsibility to set, not IT. DR is part of the BC.

Mackenzie F.: Rackspace-hosted exchange comes to mind. That burned a lot of businesses.

Christian P.: Business Continuity and Disaster Recovery are two sides of the same coin.

Scott D.: Agree, Christian.

Daniel M.: Agreed. I'm stating its not up to IT to set the BC requirements.

Scott D.: We in IT set the tone for BC by making sure DR is secure.

Natasha R.: Yes, we use Rubrik as well in the Cloud.

Daniel M.: BCP should be set, then the DR is developed to meet the BCP.

Gregory R.: Daniel, it's funny that sometimes the DR starts before the BC on the business side.

Gary C.: Totally agree. During our BC conversations with departments, they kept saying, "But isn't that IT?" The talks were very productive for everyone, to say the least.

Stephen T.: I will say that starting with Sandy, we migrated our DR environment into the Cloud deliberately. It's because we did not want to be dependent on an East Coast power grid or anything else with the international clientele. It has steadily made life easier. When we had a physical data center, it would have been both expensive and non-feasible to have a model of rebuilding an application from off-site air gapped backup as a part of our general process. Nowadays that's a straightforward thing we can aspire to get better and better at. Our model for business resilience is to get fully to the bi-coastal quarterly failover modes. Our quarterly DR test is to failover, validate, and then run in the new site for three months. At the end of those three months we do a failover test back to the original site and just go back and forth. We're not quite ready to go active. There are pricing issues. It has to do with the amount of workloads we have. But yeah, the Cloud has made disaster recovery much better. It's required the business to pay more attention and it has let us add a new DR tier, sort of between zero and one. I argue everyone should keep the number of tiers very small, because people will just invent weird things. It's simpler. But we added a new tier, which is business

processes. We need to be able to run entirely manually during an emergency. If we have an ultimate ransomware event or a takeover event, and we have to go off the internet and stay that way, we can keep the company going for a week to between two weeks to two months entirely using telephones and a few handheld devices. That's been a bit of a revelation for both management and the board.

Moderator: Okay, thank you Stephen. Ammar, do you have a follow-up comment?

Ammar Y.: Yes, regarding the tiers. I don't want to forget, because it was mentioned twice. I know that companies can go from three to four. When I joined and I introduced the idea of having three only, everyone has a different opinion on this client-facing, customer-facing. Then the corporate one that supports the client-facing and then everything corporate. This is just an idea I wanted to throw before I just jump into the main point. Cloud, as for our CEO directions we need to go in the next five years to 60% Cloud. It's not easy to achieve this goal, but we are trying. But for us it's near impossible to switch to Cloud because of the huge number of apps and regions that we have. The point is we are trying to go to Cloud, but at the same time we are expanding our data centers at the same time. It's not going to go as smooth as what we hope to because of the cost and the way that it's managed in the Cloud, and the regulations that we have. So many regulations in every continent. Every country they have their own regulations, even in Europe. In the UK we have more regulations than in Canada. It's complicated for us.

Gregory R.: We first had to unravel the DR / BC thing with the business side. There's always that conflict of interpretation. Nobody doesn't understand DR is separate from BC in some ways, and we had to unravel that. But then that was a difficulty in and marrying up our on-prem requirements with our ability to scale to the Cloud. The technology was ahead of the actual business technology, because there's a lot of on-prem devices and things like that that need to have a connection, or they have GPOs, for example, that keep them alive and keep them running. I have to unravel a lot of that business. The strategy is we have to slow down. I'm glad I saw some stuff in the chat about that. It's started with three-year migration plans and then five to six years, and we may even get to seven, because there's just things you can't unravel. We discovered it during the DR implementation when we started to build it. It's just something you don't know until you don't know.

Christian P.: To the heart of this question, if you think back in the '90s what was DR like? It was putting a whole bunch of tapes in boxes and making a whole bunch of people go across the country in order to recover systems someplace else. Then you get to the 2000s and you start having data replication going on, along with the 2010s where you've got containerization, you've got different ways of parsing your IT environment. Now that we're in the 2020s as people are moving to Cloud, obviously the strategies that you use as far as putting tapes in buckets aren't going to apply here. For the person who is saying I don't feel quite the same relationship here when it comes to us using Cloud versus other people, some of the things that have to be considered are other vendors. If you're dependent on a vendor and they're in a Cloud, what are their audiences and RPOs as far as recovering your environment? I recall one vendor where the RPO was 24 hours and the RTO was 72 hours. Then you ask somebody well, here we are on Tuesday. If your systems fail today, how comfortable are you with getting all of your systems back on Friday with Monday's data? Not many people like that kind of idea. It requires not just thinking about what you have within your own Cloud, or how you're using Infrastructure as a Service. But what about those providers that are giving you Software as a Service or Platform as a Service in the mix?

TOPIC: Separating backups from production to protect against ransomware

Moderator: Great comment, Christian. Thank you. Anyone else have a final comment on this question? Okay. What do others use to separate backups from production to protect against ransomware events?

Matthew W.: We use Veeam, Disaster Recovery as a Service (DRaaS). We offsite our backup to a Cloud Connect provider. They actually would pay extra. They move them off the storage rate. They're keeping the main backups to a different array that's air gapped. The data can only go one way, and it can't be deleted off of there, and they keep that on there for seven days, so even if somebody calls in and says someone nukes all my data out of Veeam they can't even call them and say delete it off the other side. You've got to wait seven days. We also rotate up basically it's a little NAS. It's got a bunch of 18 TB drives in it, and we actually make an additional copy to that. That gets air gapped. We have two of them. We swap that out weekly and keep that on-prem just in case we need to get data back real fast. We at least have a prior week's entire set of copies.

Moderator: Thank you for that, Matthew. Anyone else have a comment?

Gregory R.: Hi. We use Rubrik to keep the backups in the Cloud vault. Expensive, but it's nice. The immutability is a really good part of that as well. They do definitely want to ring out the licensing on the security side. There's a lot more of flexibility you get when you expand to their enterprise. But again, you have to pay for it. But as a base level, it's pretty great. I like it so far.

John K.: Is there a DR plan that is suitable to combat Ransomware?

Matthew W.: A lot of Ransomware these days is popping the DC and getting scripts into GPOs to ransom all systems, or getting into a tool like SCCM, PDQ, Intune, etc. and pushing out scripts.

Gregory R.: Direct Drive mappings as well are a vector.

TOPIC: On-site data recovery versus off-site recovery to counter ransomware threats

Moderator: This next question is by Jeff. How are companies balancing on-site data, system recovery with traditional off-site recovery effort to counter the threats of ransomware?

Jeff M.: We've talked a little bit about that. This is more of a topic probably within the last five years or so. Prior to that, I think everybody was mostly concerned with losing a site. Going to that DR site and operating out of there if you lost a site. This one's more focused on data. When you go through these exercises, what we're trying to find out is what is the balance with recovering on site in the same data center, or recovering data or systems from those restores from the air gapped copies and going through those exercises compared to how much time you spend trying to go across to the other site?

Moderator: Anyone else? Well, here's a follow-up question. What steps have companies taken to review the suitability of their DR plans for a ransomware attack?

Gregory R.: I would see that as like a part of the review of the DR plan, right? Depending on whatever interval companies would take that as. I mean how often are you testing it? How often are you revising it? I would find that's where this would take place.

Mackenzie F.: Quick question, I guess, for those with their backups and DR plans. Do many of you have, or as somebody had mentioned air gapping, having a separate NAS for that, in our case we do

have backups replicated across systems in each site. There is some segregation there. But I wouldn't say we have anything that is disconnected entirely with our backup. I'm curious if others employ that as a best practice as it stands now.

Jeff M.: From a ransomware perspective we do have what they consider storage immutability. The systems aren't on Active Directory. You can't connect to them from Active Directory. The console is not accessible through Active Directory permissions. You have to use a jump server, you have to multi-factor authenticate to access the system, and then the storage is isolated for the most part, too. We do replicate both active storage and backup storage offsite.

Mackenzie F.: Awesome. Thank you.

TOPIC: Mass communication systems used for incident/disaster notification

Moderator: All right, we'll move on. What mass communication systems are currently being used to notify stakeholders on major incidents or disasters?

Matthew W.: OnSolve is the vendor. We use them for mass communication. I think they have eight different mass communication products that they've bought over the years. We use One Call Now, or Send Word Now. I think it's Send Word Now, which is really One Call Now. But it is text, voice, pre-canned voice, text to speech, that kind of stuff, and it goes out to all users. We have a bunch of stuff pre-canned in there.

Gregory R.: For trojan / worm delivery of ransom.

Jennifer T.: We use AlertMedia for emergency notification system.

Tyler B.: Depending on the issue, we use MS Teams and / or AtHoc.

Christine H.: Each leader is in charge of notifying their teams.

Mark E.: We use InformaCast to communicate emergency events.

Daniel M.: Matthew, what was the application you mentioned?

Matthew W.: OnSolve One Call Now. They also make a few others depending on what you need to do. https://www.onsolve.com/platform-products/critical-communications/one-call-now/

Mike T.: We use AlertMedia for IT notifications as well as all other organizational notifications.

Gary C.: Old-fashioned phone trees are our last resort if all else fails. Departments are asked to keep them up to date.

Moderator: Okay, thanks Matthew. Natasha, go ahead, please.

Natasha R.: Yeah, we use xMatters, and that sends out all of our crisis management DR notifications across the whole of our business as part of that. For major incidents just internally when services go down we use our ITSM tool, which is BMC Helix.

Bryan K.: We use a product called Everbridge for our disaster notifications. That spans IT as well as just sort of operational things like inclement weather events for our staff, and also internally use things like Status Page and Opsgenie for notification to responding staff.

Gregory R.: Just to make folks feel better that don't have it, we just use Teams and good old smoke signals. Just joking. Cell phones mostly. Just to make folks feel better who don't have a cool tool.

Jennifer T.: We use AlertMedia. Similar to what someone else mentioned, you can set it up to do notification through email, text, phone call. We have a file that imports daily. That's automated to keep our user information current. It also allows us to set up different notification permissions. For example, I work in the business continuity space, and I'm a "Master Admin". But then we can set up the people like in our IT department with a notify admin and the ability to change their groups of associates that they want to notify for IT incidents.

TOPIC: Out of band communication tools

Moderator: Thank you, Jennifer. Janet has a question on what out-of-band communication tools do you use? I'm not quite sure what she means by out-of-band. Does anyone have any comment on that?

Matthew W.: Probably the same thing we were just talking about. Not email, basically.

Moderator: Okay, thank you for that. Any other comments?

Jennifer T.: One of the scenarios that we had was a cyber incident, and we were able to notify a line of business through our AlertMedia notification tool, that the bad actors and the systems are not able to view your notification. That's another nice feature.

TOPIC: Prioritizing processes and applications

Moderator: Okay, thank you. Does anyone feel they are mature in how they prioritize processes and / or applications? Would you be willing to share the detail for the buckets or criteria you've created for your organization?

Bryan K.: We find the ITSM Urgency and Impact matrix pretty useful for bucketing priorities. An example: https://www.bmc.com/blogs/impact-urgency-priority/ Not perfect, but good enough for starting conversations.

Ammar Y.: That's a tough question, because prioritization maturity is not an easy thing to do, to be honest with you. We might have some stuff, but I doubt that we have 100% maturity actually in DR, or most companies I worked with the past.

Moderator: Yeah, thanks Ammar. Jennifer, is there anything else you want to ask?

Jennifer T.: No, I just kind of wonder. I saw something, maybe an organization had four buckets, and then how they define their four buckets from a business process standpoint all the way down to where they host that application. I feel like this is something maybe if somebody has something that is in a written format of how they bucket their criteria. That would be a useful document to be able to look at.

Moderator: If anyone has that and are willing to share it, we'll make sure that that gets out to everybody. Christian?

Christian P.: Jennifer, would you differentiate that in some way between the tier ratings that I referred to earlier and something else? Or are they fairly similar?

Jennifer T.: I think it would be similar.

Christian P.: There are a few different ways that you can sort those in terms of buckets. Like somebody was saying in the chat, if you have personally identifiable information that's a part of the considerations for the tiering and prioritization. I think that has more of a relevance for the security side actually. But it is important. Some organizations will do it based upon revenue impact. If you have one server that makes $10 million versus 100 servers that make $500,000, that's what would influence tier rating. Other things can be included, such as employee impact. How many people would you have sitting on their hands if something went out versus keeping it up and running? There's also, of course, the external visibility and impact to your brand, which comes into play too. Because it may be a small system, but if it's your front door and customers then lose confidence in you, that would be a bad thing. The tiering, if you were to go out on the web and search for Gartner tiering, you'll see a few examples of that tiering that they recommend with RTOs and RPOs, as far as what fits and what categories. I like asking the business from a business impact analysis standpoint. Tell me what this would mean if you didn't have it for 48 hours, and just see how they react. Tell me, if this is down for 4 hours, is this material to you? Some organizations will take the total revenue of the entire company and divide it according to the entire year, saying if this system is down, we lose $5,000 a minute. I've been in organizations that have done that. There's different ways of weighing this and coming up with the standard. But the important thing is circulating it within your organization to get that buy-in from everybody.

Moderator: Okay, thank you. Mark? You had your hand up momentarily. Did you have a question?

Mark E.: No. I'm happy to share, and it's probably a pretty involved answer for this short time span we have. We've been doing a data governance project for the past couple of years, and what we've done there is essentially trying to answer several of the questions that were asked, and specifically to Jennifer. We went through with all of our 40 agencies here at the county, and we had them start an inventory of their applications and systems because we thought that was most meaningful to non-IT people. In that inventory we asked them a lot of questions about specific application. But then we also asked questions about DR tiering, what recovery time and recovery point they would like. We asked them to classify the data into three different data classification buckets that we made. Provided definitions for all of these so they could answer these questions, and that was our first path. Our second path this year is to go back into each of the applications and look for the data sets and inventory them. Once that's done, then we'll start diving in with the data classification and security treatments, disaster recovery treatments. Our DR planning has that information from those folks. We created a Tier 1-4, I think an hour, a day, or whatever. I don't remember. I'm happy to share that Jennifer, if you need it. But we made a Tier 0 in IT for DR planning which was basic communications, back in place phones, email, the network up at our recovery center, that kind of thing, which the average agency is not going to think about. We have our own Tier 0, and then we have 1-4. Of course they tell us what they'd like. It doesn't mean that's what they'll get. Because one hour recovery often requires a hot standby system, and we just don't have that. Anyway, that's some of our thinking, and I'd be happy to share all of that with everybody. It's been a fun project and gotten us a lot of information we needed.

Jennifer T.: Thank you.

TOPIC: DR/BC accountability and R&Rs

Moderator: Thank you very much, Mark. We greatly appreciate that. I'm going to move on to the next question. At your organization, who has the ultimate accountability for disaster recovery / business continuity and what are the roles and responsibilities? Anyone want to share?

Ammar Y.: From a DR perspective we have a DR team, from a business continuity perspective we have a business community management central, but from a tier perspective we are responsible for it. Again, like it's my responsibility. I just took over the department about 12 months ago. From a business continuity, for just T & O. Not the whole company, because we are 90,000 people. We are responsible just for the IT tech team, which is about 9,000 people. That's from a business continuity perspective. The responsibilities from a DR, like setting up all the plans, enterprise, DR, all the data centers, DR practices and exercises is our responsibility. The disaster recovery department, or we call it just disaster readiness in general. And from a business continuity, again it is just for T &O.

Jennifer T.: I have my partner in crime, Jeff, on the line here, and I feel like we have a good partnership. I'm located in the risk and compliance business area, and my focus is on business continuity. Jeff is in the IT business area, and he is responsible for disaster recovery. We call it technology recovery. I see he just dropped off. But I feel like that's worked well on the business continuity side. We usually participate in the DR exercises as an observer and engage the line of business.

Natasha R.: In our organization we have our security and resilience that are ultimately accountable for disaster recovery and business continuity. But obviously IT are responsible for delivering it. And they deliver the strategy, etc. on it. But tech will deliver. From business continuity, it's the same thing. Security and resilience will be accountable for the whole thing in its entirety. But the different business units feed in from a responsibility point of view.

Tom Q.: We are a "City within a City" with a Department of Emergency Preparedness.

John K.: We are a small shop with an immature process. Security and IT are responsible.
Niki S.: We have a Risk & Information Security Director in our IT department who is responsible for overseeing the entirety of DR / BC. He works directly with IT sub-team leads and Business Unit leads to gather relevant data and handle testing schedules.

TOPIC: Leveraging BIAs to get funding and support

Moderator: Thank you very much for that. We'll move on to some of the final questions here. This is a question on Business Impact Analysis. Is anyone leveraging those in your continuity planning to get funding and support?

Natasha R.: I would say, no. I think from our BIAs at the moment we're too immature to be able to then get funding and support from that. I think the businesses have an issue just trying to identify their BIAs on that without getting fund and support to sort of plan further remediation as part of that.

Bryan K.: https://www.gartner.com/en/information-technology/glossary/bia-business-impact-analysis

Robert G.: We didn't ask departments what they would like for RTO / RPO, we asked them what they were willing to pay for.

Moderator: Okay. Well, we're nearing the end here. Any additional questions before we sign off?

TOPIC: DR planning tools

Natasha R.: I think I've got one. From a disaster recovery point of view, what tooling do you use for your plans? Do people use a particular tooling to facilitate it? Obviously Excel spreadsheets versus I don't know, things like cutover tooling. What do people use?

Jennifer T.: We use Adaptive Business Continuity product for our business continuity recovery strategies. It aligns to Lean, Agile approach. Light on documentation, more on practice.

Moderator: Ammar?

Ammar Y.: Yes, we use Okta at the moment, and we are moving to ServiceNow, the BCM module. I used in the past LDRPS and one other one I can't remember from like 15 years ago. LDRPS became much better than when I used it. But for now we are moving to ServiceNow, which I'm not sure if it is the best, but that's the management decision to go to. Archer was a good one. If you don't have one, it's a good one now. It's Cloud-based as well. It's a SaaS and it has good features.

Moderator: Thanks, Ammar. Greg?

Gregory R.: We use good old-fashioned Excel and Word.

Moderator: High tech! All right. Well, thank you very much everyone. It's been a great discussion. We're looking forward to having all of you join us on our next Roundtable 55.

End of discussion

Products / Vendors / Technologies shared in this Roundtable 55:

Adaptive Business Continuity
AlertMedia
AtHoc
Business Continuity
Cloud
Cloud Connect
Data Center
Disaster Recovery
Disaster Recovery as a Service
DR recovery
DR site
DR solution
DR testing
DRaaS
Emergency events
Emergency notification system
InformaCast
Intune
LDRPS
MS Teams
Okta
OnSolve
PDQ
Rackspace
Ransomware
RTO/RPO
Rubrik
SCCM
Veeam
Zerto

Appendix A: All Poll Results

How often do you test/practice your DR plan:

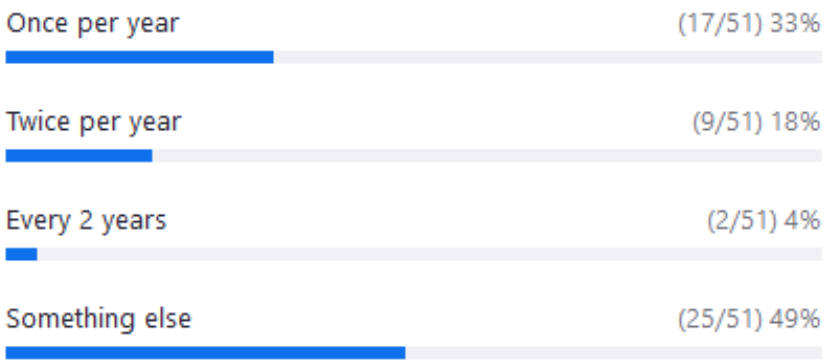

Are your DR test plans full or partial:

What process is used to determine which systems are critical for restoration:

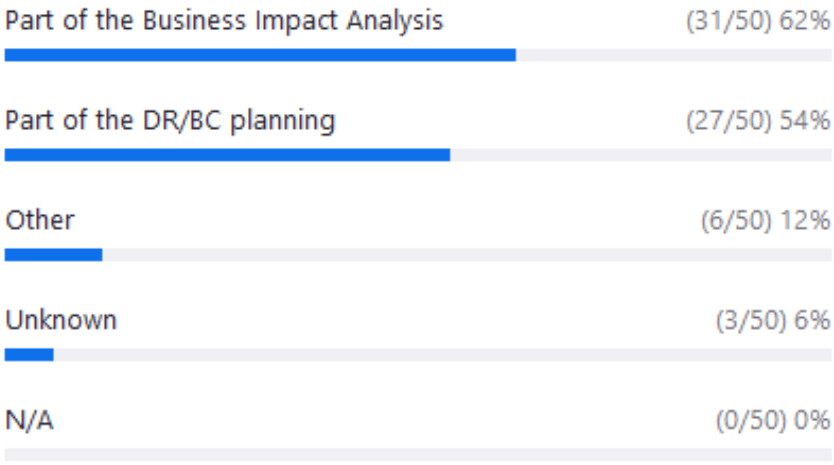

Where is your DR / BC site: