

## **DISASTER RECOVERY PLAN: POST-TALK DISCUSSION**

Following the presentation on creating a plan to recover from a disaster, there was a lively discussion and some great suggestions from everyone in attendance. The essence of this exchange is included as a .pdf with names and company references removed, in order to focus on the commentary.

## **AUDIENCE PARTICIPATION: CRITIQUE | FEEDBACK | SUGGESTIONS**

The printed 3x5 cards might be a security risk, depending on where the ring-bound card booklets are stored. The information on the cards might include contacts, contracts, passwords...

Consider where and how this information is stored.

If your company needs lots of copies, keeping these card booklets updated might be onerous.

Suggestion: Outsource this process to Oregon COOP (Continuation of Operations Plan); a subscription service. They provide a downloadable .pdf that can be transferred to a thumb drive.

The Oregon COOP might be good for large businesses or small businesses with a very small staff that can't afford time to do updates.

Suggestion for keeping information timely: Create a media wiki; an information repository. [NOTE: Check the [Rogue Tech Hub Event Calendar](#) for upcoming Tech Talk on this topic.]

## **REALITIES OF DISASTER RECOVERY PLANS**

Problem: People store things and then leave the company -or- they make a change and it doesn't get updated.

The challenge is: adding the updated information to the DRP.

The Number One Problem is: tiny, incremental changes that pile up and, over time, become a big deal.

When an employee leaves the company, a debriefing process happens; this must take into consideration outstanding issues to the plans in progress.

Debrief seems to focus on why the employee left the company so that the company doesn't replicate the process -or- make the same hiring mistake.

Disaster Recovery Team is a resource team tasked to update the wiki and make sure the wiki is backed up and stored in multiple places.

Observation: Most DRPs lack the empathy that employees are human. If there's a natural

disaster, it's easy for the CEO to announce that the staff will immediately get on a plane and fly to the company's Disaster Recovery Site ... The reality is that people are going to go home and be with their family. You need to take that into consideration when creating a DRP.

## **DISASTER RECOVERY SITES**

*Q: Does everybody here have a DR Site?*

Does 6 miles away count?

We have a fire & electrical data loss DR Site; not really an earthquake-destroyed-our-business site.

Our site is only a few miles away. The more likely disaster for us is going to be loss of power or gas leak or terrorist attack, so having it that far away isn't bad. If the disaster is big enough, it's probably going to take out more than that; the county will be at risk, and we probably won't need more.

*Q: What do you guys at the call center do if you lose phones? It's easy for us to simply pick up our cell phones and continue doing business, but what do you guys do?*

We have a physical switch that's at our DR Site that we flip and our telecom provider forwards all our calls to the new phone system. It's a big, huge red button with alarms and all, and when that switch is flipped, all the incoming calls are switched to the new center.

In the past, we've had a minor incident where someone accidentally flipped the switch ... we copied down the number and then forwarded it back to our main number. So if it's accidentally flipped, it falls over to the DR center and then is forwarded back to our site. So the plan is: in case of a real disaster, you have to flip the switch and then unforward the phones.

*Q: How many people have permission (or the duty) to go through that process? IOW, how many people do you have to educate to do that?*

Every one of our dispatchers; 40 or so.

That's brilliant! That's grass roots disaster recovery.

Yup! Everyone has a key fob with the info on it. In case of emergency, we have to figure out, logistically, who's going to be available to work. But we haven't had to do that yet.

This summer, we had a fire in our dispatch center; that was a real-world test and caused us to think about what we would do if we had to evacuate the center. One of the first things we did was to clear the brush from around the perimeter of the building so the electrical and antenna cable wiring was not threatened by burning bushes. We have a huge tower that has

a lot of radio antennas connected to it.

We also had to think about how we physically get out of the area because, in this case, they closed-off all the roads. We needed to think about how we get through the area and how we verify who we are to local law enforcement; ODOT & Police. The first thing the fire did was burn the overhead fiber to our DR Site. Back in January, we had scheduled an overhaul of our microwave connection. They were dragging their feet to complete this process. But then, during the fire, we lost the connection and our phones were down for about 30 seconds. It turned out to be OK; however, that was a heads-up moment. As it turned out, the lines were so busy that people couldn't get through anyway but it was memorable. We added that microwave back-haul so, even if everything stops working, fiber goes out, internet stops working, we're set. We have a real-time connection to our DR Site.

There are a whole bunch of regulations that require our telecom provider to notify the 911 Dispatch Center when there's a service interruption; this was initiated after some previous lawsuits. We get extra special treatment, but we pay our telecom provider a lot of money. The FCC will sue the phone companies; they have liabilities to keep the networks up and/or notify users. They sued a major telecom provider in a couple of states, and it was about a \$43 million settlement; this was due to a statewide 911 outage, specific to cell phones, and they did not report it.

I'm sure you guys know all about that.

Oh, yeah. We notify all 911 Dispatch Centers when service goes out. There are a whole bunch of ancillary text-to-911 services. [This response is from a wireless solutions architect so he lives in a world between sales & technology; supporting business-to-business sales team.]

## **5G**

*Q: What's the scoop on 5G?*

The reality is that the hype is overblown. People who buy 5G phones for \$1,000 will be beta testers. 5G is a variety of technologies. [Millimeter Wave](#) is the one that everybody is nervous about from a health perspective; however, Millimeter Wave is not a very potent signal, and that's the other challenge.

So you need a zillion sources of beams....

Yes. Many beams; it's effectively a beam-forming and beam-tracing challenge where you're following people around for miles away. But if your head is between your phone and the tower, we can't even get the signal through your head to communicate with the tower because your head will absorb the signal; it's non-reflective. Right now, we can beam through your head to get a signal to a tower, but with a Millimeter Wave, we have to find a wall to bounce the signal off of and get to the tower. It's RF, but it's really weak RF; that's why you'll have so many sites. The reality is: you don't know what you're bouncing the signal off of,

you're simply shooting it out there and hoping to get something back, and the reality is that we're hoping to send enough signals out in hopes of getting something back from your phone. So good luck making that work, at the beginning.

I'm not really interested in the extra speed, but the health issues terrify me.

Millimeter Wave is one way of doing 5G. Frankly, a lot of carriers have access to the 600 spectrum, a very low-end spectrum that does a tremendous job at penetration and does a great long-throw over distance, and we're likely to deploy a lot of our technologies in that spectrum, initially. So you get great building penetration, great blockage penetration. You'll see AT&T and Verizon will focus on Millimeter in the densified areas while we'll turn our towers out to the rural areas with a completely different spectrum.

But it will cost you more money because you need more sites and more points.

600 not so much. But Millimeter Wave, forget about it; towers are over; towers will not make millimeter work. Small cells strapped to buildings all over the place. Imagine, in downtown, one carrier to cover 5G in a downtown area, you're probably looking at 10-20 sites, which is completely different than 3 to cover the whole county. So it will be a very different world. But you'll start seeing these 2x2 metal boxes start appearing on the corners of your favorite buildings and those will be small cells replicating coverage.

They call those Micro cells?

There's really no industry standard. The interesting thing is: 5G will be a thing but we're actually starting to talk to folks about the notion of private [LTE](#) and [private LG](#).

### **PRIVATE LTE ... COMING SOON!**

A lot of folks are realizing that WiFi is not as secure a standard as a [3GPP](#) standard; it's good but it's not great. Once you get over IEEE, and you start getting into LTE standards and 5G standards, the security actually gets a lot better so this notion of talking to a municipality or a county and saying, hey we can launch small cells, you'll actually backhaul it over your own ISP so it's a non-metered connection, we'll run the network for you, you get your own custom SIM Cards to give to your folks and your own data map to administer over 4G. So now, instead of WiFi Zones and Access Points, that cost you millions of dollars to maintain, you'll have LTE Sites. You can imagine 2 sites covering two hospitals, instead of 50 data access points fighting with each other. So we're just now starting to talk to customers about their interest in that space.

Are you talking to individual users that you have relationships with or are you talking to businesses?

I'm largely focused with talking to businesses, and private LTE from a cost perspective would likely be more applicable to a business. A lot of times, the notion of looking at a municipal

network, whether that's a sensor network or a data network and then, from that perspective, assess who are the anchor tenants who would make most use of this so in some cases you would see, for example, a county government saying we'd love a private LTE network ... the cops will all use it ... all emergency services will use it. But then they'll turn around and say, gosh, we'd like to lower the cost a little bit; how can we do that? Oh, local hospital, are you willing to be a tenant and pay for your services and, effectively, we now have some LTE sites there as well?

This is the first I've heard of private LTE.

I would say that I'm all of about 6 weeks ahead of you. I was in Chicago recently and was asking a salesman about private LTE; what's it look like? His response: that's a great question. Let's see who wants to buy it and how much they're willing to pay.

And then we'll design it.

That's kind of the idea.

Let's have our marketing department design it...

Yeah.

## **FIBER SOURCES**

It's interesting that you are talking about Disaster Recovery. We talk to folks about network failover all the time; using cellular as a backup circuit. It's funny, but the first question I always get is: What's your fiber pathing? Are you using different fiber sources than we use? And, the reality is always, "Kind of. Yes." Because we have massive [ring topologies](#). While at the county level we have sites served by the same two fiber providers that you have, our devices or our towers are all connected to each other in massive multi-hundred mile rings with sources of fiber coming in from all over the place and the ability to use LTE and microwave to backhaul throughout our networks, so while we may have all of our fiber-in-common down, we can actually use tower-to-tower communication to pull fiber connections out of Washington. It's hard to even tell folks what our pathing source is. We have 68 sources of fiber in the Pacific Northwest and it all depends on how we're going to path you back out to a connection.

We had some problems with backup because we lost one of our fibers in that fire, also. And when we talk to the cellular providers, they look at how much data we use and they sort of freak-out and say that it would cost you too much money to do any backup or provide Internet for all the data you're using.

What we tend to see, and I try to be honest with customers about this, is that flipping over to a metered connection is not ideal. But, in many cases, you're in a situation where it's so business-critical ... the loss of revenue and other costs are so much greater ... that you look

at it and say, gee, if we hemorrhage \$10,000 in the next two weeks by managing this connection, what would the cost of lost business [revenue] have been? It could have been a quarter of a million dollars in lost revenue ... it *sucks* but it's a necessary evil.

What we will largely encourage our customers to do is say, "Hey, there's a difference between nominal business operations and hobble-along-because-you're-on-a-metered-connection operations." So, a lot of times, your Internet connection will know you're on failover—there's a million ways to detect that—and you can say, "Hey, now the Internet looks different." Facebook is gone. Netflix and Hulu are gone. But corporate email is up-and-running. The print servers work. Our backup system works. Our VOIP system works. But that's it. Because that's all the bandwidth we have and that's all we're willing to pay for. If you wish to use your phone to check Facebook, you're welcome to do that. But Guest WiFi shuts off. All these constraints of services are set up to ensure that we can keep mission-critical services up-and-running. We also can't afford to pay for your fun stuff. It's important to be honest about what you can afford to do. It's all about the plan; you know what you're going into and you plan for it; it's usually not as painful as you imagine.

### **DRP IN ACTION**

So, aside from the 3x5 card—the physical book—did the step-by-step process make sense to everybody? This is simply something that every business owner must do; sit down and go through what they perceive as a disaster and how that will affect their business operations. And then, simply walk through that scenario and determine how you're going to handle each task. Is there something that was overlooked in that process that you think is valuable?

You want to make sure that backups are happening regularly and that they're stored off-site and accessible. You assume that it's going to be a regional disaster and plan for that. It just has to be a geographical separation; you assume that's going to be the case and go from there.

### **FEASIBILITY OF DRP PARTNERSHIPS**

*Q: How many companies do you see that have a remote "companion" site? One local company might partner-up with a company that's in the same region and keep each other's backup; a collaboration between like businesses.*

That's a really interesting idea. I don't know of any off the top of my head, but I know that, as a company, we're very involved in associations that are important to us and we connect to other companies. To me, it would be a very natural kind of "pairing;" especially in our industry (trucking), where there's a lot of need for fixed structure; a yard, space, and having collaboration with a like-business—as an emergency partnership.

If it's somebody that you trust, and you could arrange mutually-destructive capabilities...

I think that's a very unique and cost-effective way to solve that sort of problem ... I'm thinking about other trucking companies that I've met representatives from and have all

these facilities already set up who would be in just as much of a need if they experienced a disaster, and the costs would be minimal to either side to take on additional trucks in our facility. Cost of electricity. Cost of fuel. Cost of server.

It would almost be as a courtesy.

It's a brilliant idea.

You're already seeing this kind of pairing-up in larger corporations with multiple branches; you almost have round-robin data branches ... we send our stuff East and we keep getting stuff that's sent West. Eventually, the tapes are all located somewhere else.

What about smaller companies? You're thinking huge: IBM, or other multi-national companies. In our area, we have mostly small companies.

### **CLOUD ... OR NOT?**

I think most small companies are going to use cloud-based services because that's the most cost-effective way to go ... and you don't have to rely on contracts.

You're putting all your data in someone else's facility, and there's a risk associated with that; could be fire, theft, cybersecurity ...

The companies I've looked at are doing very little cloud-based backup. They still have everything on servers in their office; they're 20 years old and full of dust. But done right, it's much cheaper.

*Q: What if the backup were stored in a safe?*

Well, with a safe, you'd have to have access to it. And you'd have to have a process to keep it updated. They do make safes with USB keys but most companies don't have something like that.

I see your point; we're kind of in this tweener zone where folks have on-premises stuff where in many cases, they're looking at moving the whole service into the cloud.

So many companies have older equipment that can't be serviced in a cloud; you can't run that stuff out there. They're scared to go that route.

### **VIRTUALIZATION**

I think, with virtualization, there's a lot of opportunity to move stuff to the cloud. As long as an older server can be virtualized, you have the opportunity to move it up there.

But nobody supports that.

One of my friends has an old [AS400](#).

Wow!

They're still viable! This is a mechanic's shop; they're not computer guys. Every time he has a problem, I suggest he go to eBay and buy a spare so he has backup parts.

They do all their payment processing, billing, ordering.

Back in the Air Force Days, with the old [RLL Hard Drives](#) that were running all the time...they took them down for a P.M. one day and these things had been running for so long that the disks were hot and, when they shut them down long enough to blow them out, they cooled off and the disks warped. So, when they powered them back on, they wouldn't boot; they just locked-up. Think about that!

Group: Agreed this was a bad situation, and not uncommon.

On the other hand, the only secure computer is one that's powered-off.

Group: Yeah. Always on or always off.

## **SECURITY**

Speaking of security, I have started reading a new book entitled, [Data and Goliath](#) by Bruce Schneier, and it's from 2015 but it seems very up-to-the-minute, and it references a website that offers updates to the information contained in the book. If you're interested in the topic of Big Data and how it's affecting privacy and security, this book is very well-written and the author seems to know his stuff. Scary!

I'll add a link to the book on the website.

## **DISASTER PREVENTION – OR – AWARENESS**

I'd say as you're thinking about Disaster Recovery, the other thing to imagine is that so many folks press down into the future and ask, "what do we do?" and they don't look back to a previous disaster and ask, "what do we do to prevent some of the things that came up?" IOW, what elements of a previous disaster might be preventable? It's a space that I keep advising customers about, with sensing; for example, I can have a message go out when the servers are getting warm. That could be nothing or it might be something to check out, like failure of an A.C. Unit. You might decide to place a box fan in the window until the A.C. Guy can come out or decide that it's a non-issue. But the number of folks who don't think through "*What can I be aware of?*" is high. For example, I have a generator ... do you monitor the fuel levels? When was that generator last started? Once a month, do you turn off power and turn on the generator to test the functionality—and make sure the fuel flow is good.

Do you have many customers using things like Tesla Wall Batteries?



Not that I've come across, but we have a lot of customers with separate propane generators to maintain resiliency. Our local community college has a whole data center run by backup generator system. All of our systems are propane generator-backed with battery backup. That tends to be the more common thing. The Power Wall thing sounds amazing but I don't know if everyone is ready to switch over to that concept.

Well, Tesla is not ready to sell to all that many people, unfortunately.

Group: Discussion about pros and cons of battery backup vs. propane generators.

We just had one of our major UPSs decide that, every time there's a power hit, it decided to turn itself off. That UPS had all of our phone switches and data switches on it.

So, what did you do?

We bought a new UPS.

We suffered downtime during the switch (to the new UPS).

*Q: How long?*

About an hour.

Oh, that's not too bad.

That was a good test because it was killing the UPS when we switched to generator test. We do weekly generator tests to make sure that everything falls over correctly and that the generator provides power, and make sure the generator's working so it will be there when we need it. Well, it was during one of these generator tests that the UPS decided to turn itself off. So the UPS is working again. We disabled the generator tests for a period of 2-3 weeks, until we had a replacement unit, and then scheduled downtime to ensure that test would be during a minimal impact time.

And, the UPS they sent us was damaged in shipment, so we had to wait for a replacement to be sent.

*Q: Where was the UPS shipped from?*

I don't know...bought it off the Internet.

*Q: How long did the replacement take to arrive?*

About a week.

We only have about 50% of our equipment on one UPS so, worse case, we could have taken the UPSs from other stations and used those in the interim.

But to do that would be a lot more downtime.

### **SENSING IMPENDING DISASTERS**

One of the things we did when we designed our building was to include an external generator transfer switch so, if a generator doesn't kick on, we can actually roll a generator up to the front of our building, hook it up, flip the transfer switch, and automatically provide power. That's one thing some people don't realize: How can you provide power to your facility if you need external power from a generator?

That's been part of our problem: we keep finding additional things that we're not monitoring. We just realized we weren't monitoring the server [IDRAC](#). The server has the IDRAC built-in, you just have to turn it on. It tells you anytime there's a power failure. But you've got to remember to turn it all on and then you've got to monitor that.

### **BALANCING RISK VS. COST REMEDIATION**

One thing to consider when creating a Disaster Recovery Plan: Balancing the risk vs. the cost of remediation. For example, one of the challenges might be: do you factor-in war with China? This would completely disconnect the supply chain from Southeast Asia. Or, do we have war in the streets of a major city that would shut down the state government. There are ways to plan business continuity around those events but sometimes, they get extremely expensive. Balancing the cost against the likelihood is difficult. One place where we struggle is: fire is pretty likely, but what's the likelihood that we will lose both of our locations? We have different ISPs so how much do you budget—or spend—to remediate those possibilities?

I might suggest that you connect with a business risk insurance agent: Tom Kaldunski at Hart Insurance is one person I happen to know; he did a CyberSecurity Seminar in February 2018; the [wrapUp](#) is posted on Rogue Tech Hub, and his contact information would be out on the website. Someone like Tom makes it their business to assess the various risks that businesses face and analyze costs based on statistical information.

That's what I'm talking about; it's not only how we remediate the disaster; it's the cost to remediate the disaster proportionate to the likelihood of the disaster happening.

You don't try to plan for specific disasters; you plan for general things happening, and then your cost analysis is what is the cost of *a* disaster? You might decide that it's not cost-effective to do anything; data lost.

### **COMMUNITY**

What other groups do you associate with that have the same types of risks that your

company has? I'm talking about partnerships. If only talking about what each other does.

In the 911 World, it's competitive because the state wants us to consolidate and we don't know what that would mean. Would that mean downsizing from forty-four 911 Call Centers to 10? Would we be one of the 10 remaining? So we try to be the best, in case something like that does happen because the likelihood of consolidation is still high, but we have the best center in our area and certainly, regionally. If there is consolidation, it's going to happen regionally and we have one of the better centers.

### **STANDARD PRACTICES**

Even getting to the consolidation...I'm out talking to the agencies that the Dispatch Center supports and because the center's right here, they're largely unified but, as I head out throughout the state, I realize that everybody uses different software ... nobody talks to each other ... it's incredible how uncoordinated emergency response is ... I wouldn't even say this is confined to our state; it's across the country. There's no gold standard ... this is cheaper ... this is good enough ... we've always done it this way ... I don't see the benefit in new software ...

Is that, in part, because there's so much attachment to government in that particular realm?

I think so. The people who are in that position have typically been there for 20 years; it's not like it's a modern, evolving technology.

And, what's their background? They probably don't have quite the background in technology.

Absolutely. Most of them have come up through the ranks, after being a dispatcher or a records clerk. And now, they're running the center. It's always been like that.

So, their thought process is different.

At the agency level, strangely, you're actually looking, in many cases, at the lowest common denominator. You want to keep your folks out doing their job, and simpler systems, even though they may be less-effective systems, are systems that people are comfortable with, especially when your workforce has an average of 15 years of experience. They're great cops but horrible I.T. Guys. So, the system looks like DOS, but they can run it. And, honestly, that's one of the biggest barriers; these guys know how to run it and, if I give them something new, a third of them will probably take early retirement.

They don't want to have to deal with failure.

It's no country for old men.

### **LEARNING NEW STUFF...**

It's the same thing in Field Services where it's clear that so many businesses have to move to

wireless work orders and mobile payment processing. I'll tell them that I have a basket of technologies that are vastly cheaper than, frankly, the press-hard-you're-making-three-copies paperwork, let alone the efficiency, and they'll say yeah, but my HVAC Guy has been doing this for 28 years and knows everything about HVAC and ...

And he's got a whole carton of forms in his truck...

Exactly. What do you want me to do? Fire my best technician because he has to learn something new? The seasoned workforce is intensely loyal but rarely want to learn new stuff.

I don't know that it's necessarily law enforcement or whatever. Drivers are the same way.

My goodness! As we went around trying to sell [ELD \(Electronic Logging Devices\)](#) as the mandate, the number of folks who would say, I can't do this. And I would reply, "it's a legal requirement!"

Even if your drivers can't handle it, you must do this! This requirement has been in the pipeline for 15 years. We can't trust drivers to log their hours on a piece of paper. This is mandatory for all truck drivers to log their hours electronically using an ELD.

Owner/Operators are the only ones who have been saved at this point.

This triggered a discussion about ELD and whether or not it will actually happen.

Need a machine that interfaces with the driver's vehicle that records when the vehicle is moving down the road. The writing was on the wall a decade ago.

And it had already been delayed, and the prevalent opinion was that it would not be put in effect.

People who were our I.T. Representatives accused trucking company owners of drinking the Kool-Aid and said the law would never happen. Well...here we are!

New Tech: [Hololens 2](#)

That's what's amazing; what we noticed when we switched is: it's better. It's vastly better from an efficiency perspective, from an accuracy perspective. It's always been the case that we've considered putting computers in our trucks for tracking and messaging and that kind of stuff but, as soon as we started going to [ELOG](#), we have drivers who are anxious to have the company repair their devices quickly because they don't want to go back to paper logs.

Drivers have complained about the fact that they have to inspect the vehicle. In a scenario where you have a reefer trailer behind the cab, the driver has to monitor temperature by looking in the mirror at the dial to see what's going on. And there are folks who are doing

remote telemetry to log where their driver is and the temperature of the load. They can see when the load is getting warm and actually call ahead to find a refrigerated location and request the driver to pull off at a specific mile marker exit so the driver can do an off-load and reload of the cargo (e.g., steaks) in a cooler trailer. While this will take up a lot of time, it's better than the driver showing up at the destination with a "dead" load—and not get paid.

One of our customers does high-dollar wine shipments, so the temperature control is extremely important. So, for us, we have all that telemetry on our trailers; we can see if we are one degree out-of-spec. Driver must check if the source of the temperature issue is the trailer or something else.

One person shared some information about a local vineyard Wines: they deliberately subject their wines to hot and freezing temperatures ("skunk" their wines) in order to harden it because they don't want their customers (who put the purchased bottles of wine in the trunk of their car) to get their wine home and discover it has turned to vinegar. By passing their wine through extreme temperatures, it alters the wine so it's hardened against turning to vinegar.

Another person shared a story about some of their customers ordering one whole reefer [exclusively] for two cases of wine in order to ensure that the temperature is controlled. The wine is obviously very expensive and the shipping costs are a small factor in the overall delivery process. So, if you log the trip, you have proof that the temperature of the truck was constant.

### **UPCOMING MANDATES -OR- REGULATION**

You're ahead of the curve. FDA and USDA are imposing supply chain liability either the end of this year or early next year, as far as temperature monitoring goes. So, from the moment the supplier begins producing your food all the way until it's delivered to the grocery store, the temperature must be tracked, and stay within the safe zone the whole way. Everyone has to log it and track it.

We in the supply chain actually have liability.

Exactly. So they work their way backwards from the point of sale. If your product makes a bunch of people sick, they start with the grocery store and request temperature logs and, if the grocery store doesn't have those records, they don't go any further; the grocery store is liable. Whoever, in the supply chain, is unable to prove that they monitored temperature, is liable.

Because of the risk of this type of investigation, we're going to make sure that everything is done right.

I'm guessing that a lot of your competition is going to evaporate.

It's easy to forget the cost of transportation of food. We have done a lot, over the past 20 years, to make the cost of transportation higher and higher. So, stuff like this is only going to increase it even more—on foods, goods, everything. One of our specialties is that we do a lot of specialized loads; oversized and other specialty loads (like large sheets of glass and heavy equipment), so that will increase costs. When you have a job site that requires specialized equipment, the transportation costs will be greater; That increases the cost of construction. Building costs get more expensive. Rent gets more expensive. The transportation chain can really elevate those kinds of costs.

*Q: For better or worse ... why do we want ELOGs?*

Well, because we don't want drivers falling asleep at the wheel of an 80,000-pound death machine and ramming into a busload of nuns.

As they add more regulation, the smaller businesses will dry-up and you're going to see that Amazon Trucking is the only business that can afford trucking.

Amazon is one of the largest trucking companies in the country because everyone is struggling with supply chain issues.

### **MOVING TOWARDS VERTICAL INTEGRATION**

It's amazing because Amazon is now the third-largest trucking company in the country, behind UPS and FedEx, and they're the 4th or 5th-largest cargo airline in the United States. This is because of the incredible demand for supply chain logistics tracking. They've decided to do it themselves because they must have control. But that gives you an idea of how big Amazon is, with vertically-integrated trucking and airline industries.

I read an article last week, by a financial analyst, that stated this is a great time for Amazon to consider buying FedEx...! They need that much logistics control and the stock price is just right to consider this move.

That triggered a general discussion about Amazon's need—and ability—to acquire FedEx. Consider: Amazon today is like Sears of the past century ... and look what's happening to Sears now; they're going OOB.

I saw today that he (Eddie Lampert) got his buyout approved. I also saw that he's simply breaking up the company and finding ways around the (pension) obligations.

That triggered a general discussion about retirement and who will pay pension obligations ... taxpayers.

The service sector will soon be automated enough that those jobs will no longer be available to retiring workers, as in the past.

Amazon is hiring retirees who travel the country in their motor homes and park in Amazon facility parking lots; they're slower than robots but they are better workers. Amazon actually has a policy that prefers retirees.

That triggered a general discussion about where Amazon is going...mobile fab lab; assembling stuff (like cars and other items) while on the road.

### **AMAZON**

Amazon's website is perfect! It always works, their customer service is excellent and there's rarely an issue they don't promptly resolve. Everyone agreed.

Talk about an interesting industry: think about space travel. Bezos is sinking about a billion dollars a year into space travel, with zero dollars in revenue.

That triggered a group discussion on space travel...as soon as somebody is making money at that, others will join-in.

### **MOVING TO MARS**

We keep laughing about Elon Musk and some of the projects he has invested in. He'll never admit it, but it's obvious that his tunnel-boring empire is everything to do with Mars; planning to send a tunnel-boring robot to Mars and build an underground habitat.

That triggered a general discussion about how human habitation will look in the future. Health issues, logistics, getting stuff to grow on Mars, becoming self-sustaining.

It's hard to think in the time scale that would be necessary.

*Q: What's the length of time it takes for a one-way trip to Mars?*

4-5 months.

*Q: Does anyone else watch The Expanse?*

Oh, yes! Very well-done.

And, they pay attention to actual physics!

Yes!

They're already discovering massive logistical problems for humans: with, for example, cerebral spinal fluid building up in the brain and actually causing brain damage...around the 6-month mark. There is concern that, without simulated gravity, it's possible that someone could turn into a vegetable.

If you go to Mars? And you're there 6 months?

Any micro-gravity environment. We're built to have our cerebral spinal fluid drain, following gravity and without gravity, that fluid builds up around your brain. You have to start physically putting people in centrifuges to pull the fluid out.

Bones need it to stay strong.

And, with *The Expanse*, you see what starts happening when these constraints play out evolutionarily, you start seeing different physical forms of people.

I saw something about an astronaut who came down from the space station after 200 or so days, he could barely walk; he had to basically re-learn how to walk.

Yes. Some people say that those people simply lost their balance. No. Quite literally, he has a form of motor neuron brain damage. He has to wait until that brain fluid drains out.

There's changes to vision and other senses.

Even changes to DNA, because they sent one twin up and checked both twins afterwards and noticed differences.

Scott Kelly and Mark Kelly.

You can't just hang upside down for a while and reverse that?

No. You need something to pull you. You'd simply put yourself in a centrifuge with your feet at the bottom and spin around to push all the fluids down to your feet.

Back in the 90s I guessed completely wrong about Amazon; I thought that Amazon was going to lose because books, which was the only thing that Amazon was doing (at the start), were going to be print-on-demand stuff, and that was going to be the solution. Local bookstores would have relatively inexpensive full presses with high-quality printing, book covers.

And coffee while you wait for your book to be printed.

### **3D PRINTERS**

Q: Speaking of 3D printers, have you seen the [Replicator](#)? This printing process works like a reverse CT scan and uses 2D photos of an existing model to cast images onto a cylinder filled with a synthetic resin called acrylate. The key is that the resulting object can encase existing objects in resin.



I noticed that one of the techniques for artificial meat that is actually made from real cells, is made via a 3D printing process. No animals harmed.

So you order a hamburger from Amazon and print it in your house.

They're working on human organs.

Next stop: Order your 3D printed girlfriend.

Kind of already happening now...with robotics.

### **REQUEST FOR TOPICS (RFT)**

We need some topics for the upcoming meetings. March 7th, we'll be here again. What kind of topics do you want to hear...and, if you mention something, it might be you presenting.

### **INTRA-WIKI**

Internal data communication. I've been thinking about the wiki that I put together and how that might be a good way for companies to pass data between departments, between people, and store and update data. I know a lot of companies share word docs on shared drives.

Skype.

Slack.

So, an internal wiki. An iWiki? No? What would you call that?

An Internal Wiki.

I was shortening it for people who are syllabically-challenged.

Or, you can call it an intranet, but that's less-precise.

Knowledge Management.

Can I volunteer someone else for a Tech Talk? I'd call it 911: All The stuff you didn't know but you really want to know.

What do you know? We already had that Tech Talk. The [wrapUp](#) is posted on RogueTechHub.

And, he did a really great job; he's a good presenter!

I'll have to read it.

It's actually a very good read; Karen did an excellent job of transcribing the talk.

I'd love to have both of you do Tech Talks; you're both great presenters!

I think the thing that is most fascinating to me is: *How do you manage the government bureaucracy that peers in?* Because I've done some I.T. Consulting for a defense sub-contractor and the government weight—the D.O.D.—that presses in on that is insane. The regulatory requirements are ridiculous for a small operation—we're talking like a 20-man shop—that are being treated like they're a Boeing or a Lockheed Martin. Their expectations are insane and I can only imagine that you guys have to deal with similar issues.

It's not as bad as you think because the technology behind 911 is ancient so we're still running the same ancient regulatory processes that we have been running since the 60s. Really. There is, obviously, the security aspect—the FBI Aspect—when it comes to data; the [CJIS](#) aspect to it. But, operationally, it's kind of like the Wild Wild West between all the agencies. You really have no clearly-defined structured processes for receiving calls. Now there are obviously guidelines that you want to follow, but when you start talking about official processes like using software to control part of your business process, it gets super expensive. We have smaller centers; for example, in another part of the state, we only have one or two dispatchers that work by themselves, and they are also the janitors and perform other in-house duties. They answer the phone and say, "Hey, BillyBob, there's a call for you; run over there..." Calls are still recorded on Word documents with word-of-mouth. There may be a cell phone here and there but ...

Then you have huge centers like in a major city where their emergency communications center has multiple sites connected together. You have so much bureaucracy to try to get anything done ... it's so slow-moving that they are still doing things that should have been done ten years ago ... they're still trying to get everybody to sign-off on that. So, locally, we are in a unique position; we're small enough that we can change and keep up with technology needs but when you start getting big, it gets difficult because of the politics involved. There's not so much regulatory pressure from the government as you would think with 911. Obviously, we have to answer the calls, we have to perform other procedures. We have to conform to CJIS Policies. I use it like a hammer on my dispatchers, often reminding them that (for example) they cannot stay logged-in.

So, what's the title for that talk?

[Laughter]

I'll come up with something.

Who wants to talk on March 7th?

I might have a conflict.

Aren't you a principal at your company? Can't you simply announce your obligations?

I might have a personal conflict.

## **DEALING WITH EOL SOFTWARE**

Another topic I've been thinking about is: Dealing with End Of Life Software. A big example is Windows 7; that reaches EOL in January. It's ridiculous for us to pay for continued support from Microsoft. So we're coming up with: How do we upgrade? What does that upgrade path look like? What decisions go into that upgrade?

That's good. EOL Software.

Yes. And there's a lot of software that goes EOL over the next two years. You've got CentOS 6, Red Hat 6, which is a major issue because of Linux; there's been a huge explosion in the business sector using Linux, especially smaller businesses using Linux, happened about 10 years ago. And many businesses have to make a decision soon.

That's a big change because then you have to move to [systemd](#); that's a huge change.

Yes. And 8 is almost out. But 8's not out yet. So you have to consider whether to move to 7 or wait. Server 2008, domain controllers ... lots of EOL software that's coming up in the next two years. So if somebody has some experience going through that transition, and wants to share their migration experience... I will have that experience in the next two years, whether I want to or not.

What do you *imagine* it's going to be like?

I imagine it's going to be very painful.

What if you do a Tech Talk on Imagining EOL Software?

Then I would want to imagine unicorns and sunshine. And I think the reality is going to be a lot more pain, user suffering and retraining. Because, on the server side, now you're dealing with application changes, and changing applications does that involve changing user interfaces. If you're changing to Windows 10 ... I think that's where I'd be most interested because I know that companies went through this with Windows XP and that was never really a problem.

## **UPGRADING WINDOWS**

Windows 10 has a compatibility problem.

I'm not terribly worried about our software working on Windows 10 because we do have *some* Windows 10 machines in the company; like 3-5 Windows 10 machines that have broken

and been replaced. And all the other of our 120 machines are Windows 7. So, how do we upgrade in a way that is user friendly? That's a big issue. We have a really lean I.T. Team and our number one problem is frivolous user complaints. We try to be lean on the user side because we want to maintain the server side and the network side. So, if we upgraded to Windows 10, our users might say, this doesn't work because of your stupid Windows 10 and that's the problem, when the solution is to simply click the button.

I could certainly do a talk on that because I've probably done 500 Windows 7 to Windows 10 upgrades.

You have?

Yes.

We have a speaker!

We're in the process of ... we have about 75 left out of all of our group, and we're trying to do it before January (2020). So, we're mandated by the Feds to have all the upgrades completed by January 2020. One of the challenges we have is that our users are field units, so we don't see those units in our shop. We have to schedule upgrades with our users while they're trying to use some of these vehicles 24/7. So, if we have a machine on a fire engine, we can't take the fire engine out of service so we can swap-out their software.

That's a logistics challenge!

But we have some workarounds that we've done to help with that. And, the bigger problem with that, in my opinion, is that it changes every 6 months. So, when you're upgrading Windows 10, which version are you upgrading to; is it 18.09? Is it 18.03? Or, guess what? You when you just finished upgrading to 18.09, the next upgrade comes out in April and now you have to upgrade to that.

And the thing is: it wants to reboot on some of them and you can't do anything about that, unless you have the Enterprise Edition and you're running your own [WSUS \(Windows Server Update Services\)](#); that's really expensive.

You can get help with [GPOs \(Group Policy Objects\)](#) and stuff.

The key is to run your own WSUS server.

Yes, and control that. Because I can't have my dispatchers restarting in the middle of the night when they're in the middle of a call. There's no way that we would let them restart. So we control all that. They get notifications and have a 4-hour window to restart, so they can control that. It's not just Enterprise. You can control that. Now in some cases, we don't push out the updates at all. We don't do the large upgrades of Windows from one version to the

next.

Those have been problematic.

We're developing a process right now, and we've almost got it developed, where we can actually get a clean install of Windows and deploy all of our applications in less than 15 minutes. So that's our goal right now. We're not doing a re-imaging; we're doing a clean install and a deployment of all the software in 15 minutes so no matter which version of Windows is out, we can upgrade it, deploy it, and make sure it's going to work for us.

How feasible is for you to use something like [VMware](#) or [Xen](#) or [KVM](#)?

The problem with virtualization is: the dispatchers have six monitors, and virtualization with that many monitors requires so much backend.

Ah, yes; you have to do [VDI \(Virtual Desktop Infrastructure\)](#).

Yes. We've actually tested VDI; it allows you to draw all the stuff on your screens here and have all the servers in the back room. It's a performance issue. For us, performance is number one, so we don't want a single delay. If we have a 3-second delay trying to query an address, it's an issue. We are monitored; every single second, we have a metric we have to meet. Our infrastructure, our servers—our clients—all have dual SSDs in them, run in a RAID configuration for performance; they're mirrored. So we don't get the same level of performance.

It's redundancy; not a proof of performance.

No. No. It's [RAID 0; striped](#).

RAID 0 is mirrored.

Then RAID 1; striped.

[\[RAID 0 data is striped \(split\) evenly across all disks in the RAID 0 setup. RAID 1 data is mirrored; data is fully stored on each disk.\]](#)

You know who was talking about that issue was at the CORE Tech Talk they had recently; the techs were all talking about exactly the kinds of things you're talking about and they were all pissing and moaning about what a pain it was to go through this process. Anybody who had successfully achieved an upgrade was lauded by all.

Yeah, it's a challenge, depending on how many units you have. I think about someone like a large hospital; they have about 2,000 workstations they're trying to upgrade; they just finished the Windows 7 upgrade, and now they're starting the Windows 10 upgrade. And, for

them, compliancy is a huge issue, with HIPA Compliancy and all those other things they have to maintain. Fortunately, I think they're going more virtualized with a lot of their products—their EPIC products—and that will make things easier.

I'll make a suggestion. What we did at [a military facility] was to basically revoke all privileges for users. They got access to certain things they needed to do and nothing else. So we were able to manage six major commands and minor commands which were geographically-separated, with 100,000 users, tens of thousands of machines, all managed by five people in Virginia and three guys at the major sites. That's it. Everything ran. And, it was always up. And this is mission-critical stuff. Now, of course, this is military so it was real easy to take away privileges from these users. People didn't like that but it worked.

One of the things that we found is that ... we've been talking about restricting user access, too. Our department deals with all 48 [contiguous] states and a few Canadian provinces, as far as applying for permits. And, all the states have pretty much switched-over to a web application; there's only around six different companies. But each state has their own unique section of that and for a long time they were still running very old web software on their servers. So we were having lots of problems with Java because they only wanted Java 5 and then they went to Java 6, and one state requires Java 6 and another only permits Java 7 and the users are trying to figure out how to access the various versions of Java-based mapping software in the various states.

## **VIRTUALIZATION**

In the Unix World, there are ways to deal with that. In the high-performance center that I work with, as you can imagine, we've got scientists who prefer Python 5.75 and others who prefer Python 5.76, and each develops dependencies on their respective version. We have to have every one of those environments available for their particular programming and not interfere with others. So when they go to run their job, they have to set up their cluster with their preferred environment, pointing to the correct version of Java, Python, Ruby, whatever they're using. It's actually fairly simple. There's a command in Unix called Module that's specifically designed to accommodate this situation. In Windows...

Virtualization ... you virtualize different environments, like Docker does.

That's a really interesting idea.

## **AWS LAMBDA FUNCTION**

Especially if you go with AWS because AWS is now doing layers with the [Lambda Function](#). Lambda Function is serverless ... basically, you write a web API and, instead of standing up a server and then making it run as a restful request underneath Apache, it simply runs standalone in a micro container. So there's no server to support. There's no nothing. All you do is write a little bit of code and that's the Lambda Function. Anytime someone requests this URL, you pop it up and execute it and every single person around the world who makes that request gets their own separate instance of that, so it's infinitely-scalable with no

performance hit.

The problem is: if you're doing something like a PHP Function ... I'm writing this in PHP ... then I need the right version of PHP and I need the set of libraries it relies upon, and all that makes the Lambda Function huge. But what they've done now is: layers.

But when you create layers, you can have PHP in one layer and build additional layers depending on what you need. OK. PHP 5.7; that's my layer; there's one common layer for that; then I simply write my little function and...

... build on top of the different layers, depending on what you need. I see. Interesting.

And now, they're actually moving that technology into the data centers so you can get AWS Cloud Services deployed locally.

AWS is pretty amazing.

They are astounding!

Yeah. And they're getting better every day.

Tomorrow night. That Amazon Consumerism video. You make all the money on the shipping side or make all the money on the AWS side.

It's getting cheaper, though. It's not as expensive as it used to be. I'm sure that, at some point, I'm sure we'll switch over to it.

## **JUSTIFYING CLOUD SERVICES**

I would love to see us move to Cloud Services but when you start to price stuff out, it's thousands of dollars per year, and it's hard to justify spending thousands of dollars per year versus I can go buy a server from Dell for \$5,000.

Yes. But what's your labor cost ... what's your downtime cost while you're doing the transition?

Yes. There are a whole bunch of soft costs that need to be factored-in.

Just as an example, I was looking at some of the AWS Pricing today, as an actual customer. I set up a wiki in an Amazon Cloud; it's [dual-homed](#), so it's redundant. It's using their Aurora Database which is [Postgres- and MySQL-compatible](#), and it automatically scales plus it's automatically separated across multiple physical zones so it's highly-available and highly-redundant, and then it's fronted by a load balancer that distributes it across the instances that are running. But the cost to run that thing is going to be about five dollars per month.

Five dollars?

Per what?

Per month. For a wiki that's fully redundant with separate databases.

But doesn't that depend on how many users are running it?

Oh, absolutely.

But if you scale it out then ...

It's separate charges: How many tens of thousands of requests are made against the database; that's what you're charged by. It's something like a tenth-of-a-cent for every 10,000 I/O requests to the database. And there are similar charges for other servers, depending on the instance size you're running. I used a [t2 Micro](#), which is really low-end service, but if I need more performance, I can set up multiple t2 Micros or go to a larger one that costs more.

Still, five dollars per month is not that bad.

That's basically free. So, How To Use AWS might be a good Tech Talk topic. I don't know how to use AWS.

You missed the Tech Talk on AWS about 4 years ago, but the wrapUp is posted on Rogue Tech Hub.

*Q: Want to do a Tech Talk on AWS?*

Yes. OK.

I'm thinking about what we do with our internal wiki and, if I could get it in the cloud, highly-redundant, for five dollars a month, I'd do that in a heartbeat!

OK.

Because, that means that, whatever happens to our physical location, I have information that can be stored in a way that's easily-accessible, highly-redundant, highly-available...

Well, one of the reasons it's so cheap for me is that I only set up one geographic location. I could automatically replicate that anywhere in the world, throughout their availability zones and geographic regions. So, there are 8-12 regions. US East One is Northern Virginia. East Two is in Iowa. US West is Northern Oregon. And then, each of those regions has multiple availability zones, which are physically separated far enough that they're supposed to be very



resistant to natural disasters in the same zones. So when I set this up, I set it up in East One, in Northern Virginia, because that's where the FDA is, so it's only in that region, but that region has five availability zones spread around the D.C. Area, so it's highly-redundant. But I could go an extra step and say I want it replicated in the Portland Area. One of the reasons you'd do that is something called [CloudFront](#). Are you familiar with [Akamai](#)?

## **CLOUD SERVICES**

I can spell it but I don't know what it is.

It's a Content Delivery Network. Basically, they buffer—or cache—large amounts of the Internet in multiple locations; that's what CloudFront is, in front of AWS, sitting on top of their [AWS'] pipelines. So if you set up a website [locally] but a lot of your users are in Australia, and your servers are in the U.S., you set up a CloudFront Service in Australia and cache most of your data locally so the web response is faster.

I would be very interested in that, myself.

OK.

You're talking about stuff that ... my best understanding of how to use AWS effectively, you've just cut the cost by a factor of ten. And that was me assuming I was doing something barely bare-bones. So if I can barely bare-bones the whole thing at five bucks a month, when I was thinking I was going to be spending fifty bucks a month, that's revolutionary!

Amazon also has something called Free Tier; that in D2 Micro. I run business websites in Free Tier; absolutely free.

So, March 7th or March 14th?

[sigh]

March 7th gives you a whole month. March 14th gives you 40 days.

I'm just dipping my toe in ...

Your big toe is far better than my completely undipped toe. It's one of those things that, if I could just get started with exactly what you described, that would be great!

Could you just bring a laptop and walk us through spinning up a machine? We'd love it! That would be fun!

If you could describe exactly the setup you just described, in detail, and how I could do it, that would be great!

Let's do it on the 14th because we can go to Luna; they have a better presentation room.

OK.

Then you can present on the 7th.

OK.

More group discussion on various aspects of AWS, BackBlaze, Glacier products as backup.

<http://roguetechhub.com/wp-content/uploads/2019/03/RogueTechiesDRPWrapUp20190207Discussion.pdf>