Strength in numbers: How networks enhance DER management systems

Metcalfe's law states that the value of a telecommunications network is proportional to the square of the number of connected users of the system (n^2).

Anyone who regularly uses LinkedIn for sales, recruiting or job hunting knows the value of your network is directly proportional to the size of your network. After all, if you only had one or two contacts on LinkedIn, you wouldn't have access to all that many second- and third-level contacts, would you?

This correlation between network size and network value reflects Metcalfe's Law, which states that the value of a network corresponds to the square of the number of nodes within the network itself.

Given the proliferation of behind-the-meter solar systems and battery energy storage, Metcalfe's Law is one reason power industry players should be thinking of ways to leverage distributed energy resources (DERs). That's because intelligent management of DERs can go way beyond traditional load shedding or demand response. With a big enough network and bi-directional communications, DERs can deliver continuous and precise capacity that supports the grid in a variety of ways, including curtailment and spinning reserves, as well as voltage and frequency regulation.

In fact, through the power of a network, DERs can provide gains to several power-sector participants. That includes utilities, system operators, energy service providers (ESPs) and customers themselves. Here's a quick look at how networks deliver benefit.

A set-up for value

The interrelationship between network size and value was first presented in reference to telecommunications networks and later applied to the Internet by Robert Metcalfe, an electrical engineer who was one of the inventors of the Ethernet.

As this illustration from Wikipedia shows, two telephones can only connect to each other, but five phones can make 10 connections, and twelve phones can create 66 connections. That's one way to look at Metcalfe's Law. Another is by considering the nodes of the network in terms of sets.

If you have three items, you have three sets of one, two sets of two and one set of three, which gives you a total of six sets. When you have four items, you have four sets of one, six sets that contain two, three sets of three and one set of four. That gives you 15 sets. As the number of items goes up, the number of permutations of sets that you have creates a rapidly rising exponential curve.
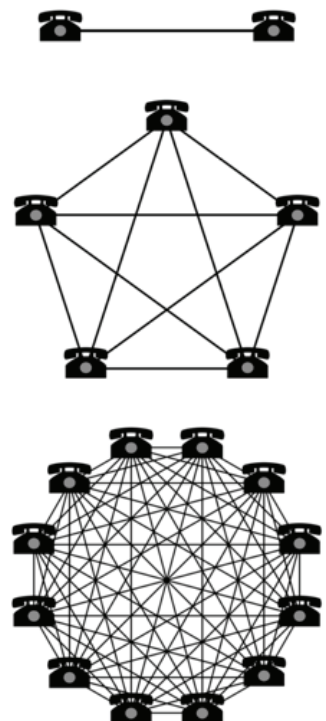
It's this approach – viewing network nodes as sets of items – that enables maximum flexibility and capability when leveraging distributed energy resources to provide grid services to a utility or independent system operator (ISO).

Why? Because viewing DERs within the construct of sets allows you to arrange which grouping of assets you'll engage to deliver the schedule and controllability you require for any given service. Generally, you won't be able to use all your DER assets simultaneously, but you will be able to use some of them at any given moment. You may even be able to use fractional amounts of the process storage or energy a DER asset has to contribute to grid support. If you have 10 megawatts of resources connected, you generally won't get 10 megawatts of capacity 100 percent of the time, but you'll probably get two or three megawatts on a constant basis.

What Metcalfe's law does is allow you to get closer to that 10 megawatts of potential because, at any given time, some of the connected resources will be available to move up in power consumption, while some will be available to move down. For instance, if you have a building that's too hot, you can turn the air conditioning system on but you can't turn it off. On the other hand, you might have a connected load that's filling the reservoir of a drinking-water system that's nearly full, so the water utility team would be more than happy to have you shut off the motors on the reservoir's pump. Add storage to your DER mix, and your potential for grid support services expands even further.

Meanwhile, multiple loads, storage devices or PV installations can appear to a utility or ISO/RTO as one resource with plenty of capacity rather than many resources with limited capacity. Because of your ability to group assets – and because Metcalfe's law ensures that there are many potential groupings – you're more likely to get close to that 10 megawatts of potential.

Right on schedule

Along with more reliable capacity, what other benefit does Metcalfe's Law bring? For one thing, it eliminates the scheduling barriers to participation that so many traditional demand response programs impose.

Right now, there is often a disconnection between the scheduling intervals that a utility or system operator would like to use and the scheduling intervals that a load or another resource uses.

For example, Enbala supports market operations at a major independent system operator that runs its market on hourly intervals. There isn't some inherent reason for that interval in a physics sense, but it does make for an easy-to-run market, one that provides enough time for participants to submit bids for things like regulation service, as well as time for the system operator to decide what it is going to do with those bids. This market runs from the top of the hour to the top of the next hour.

Now, suppose you have a large water pump that could potentially participate in some kind of grid service, but its production schedule goes from 7:30 a.m. to 5:30 p.m. and it has two 15-minute periods during the day where operators are on break, presenting at least four hours during the day when participating in the aforementioned market wouldn't work. If each of the resources in this system operator's territory had similar misalignments, the number of potential participants drops substantially.

Such misalignments can easily result from one-to-one relationships. If Pump A is to participate in Market B, their schedules have to line up perfectly, because with a one-to-one relationship, the DER is either in the market or it's out. There's no fractional participation. Through a network effect, assets are able to fractionally participate, not just in terms of capacity as described previously but also in terms of time as well. Because of the flexibility networks present, odds are that when one player is fractionally out, another is fractionally in.

Here, again, as far as the utility or market operator is concerned, all those DERs appear as one stable resource provided by an aggregator. From the customers' points of view, this provides an opportunity to maintain production schedules and site comfort levels without having to adjust their own schedules and align themselves to the market operator or utility's schedule.

Under control

Metcalfe's Law also makes a collection of distributed energy resources more dispatchable because, for the most part, grid services are most useful if they're the equivalent of a continuously tune-able knob. So, a utility, system
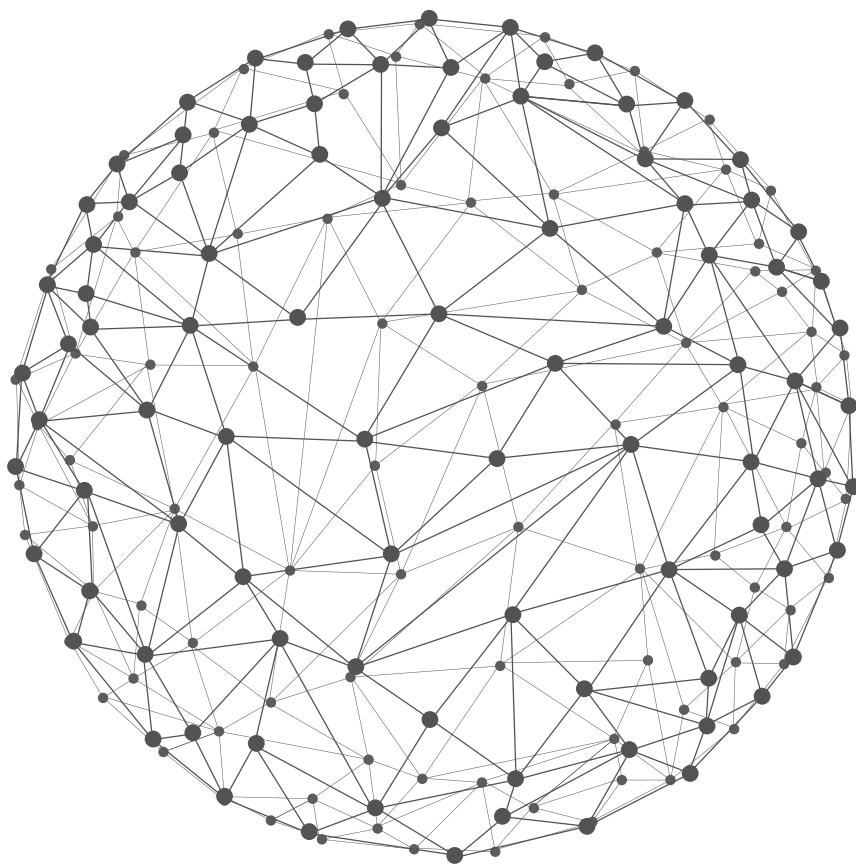
operator or ESP might say, "Right now, we need 6.93572 megawatts." A few seconds later, they need 0.34875 more, and then 1.86 less. With a network, the utility or grid operator can dictate how much capacity they need with whatever precision they choose.

This is true even when many of the resources are not, in themselves, continuously adjustable. There are many resources you just turn on and off. And, without a network aggregator like Enbala, such resources might be excluded from participating in ancillary services because they don't have the control characteristics the utility needs.

But, if you are able to take advantage of the exponentially growing numbers of combinations of sets within a network, you can mix resources in a vast number of permutations to immediately deliver exactly what the power supplier wanted. There is always some grouping of DERs that will deliver within the exact parameters and with the degree of precision the utility, ESP or grid operator requires.

And, networks can be leveraged by an organization's trading desk, too. Traders get the same information dispatchers get, allowing them to optimize fuel stack or buy-sell contracts.

Harnessing the power of DER networks

Because of Metcalfe's Law, networks are inherently useful and flexible. For a collection of DERs to deliver top-tier value, a number of characteristics should come into play.

For one thing, the aggregator's software platform should be a continuously operational, real-time solution so that power providers can leverage one system to enable multiple applications, such as regulation service, Volt/VAR management, fast demand response, renewable firming and contingency reserve. Achieving multiple grid services with one infrastructure leads to lower cost of ownership for the system overall.

To achieve that broad scope of functionality, the platform needs to have distributed computing nodes that integrate with DERs, execute local optimization routines and relay state information. That way the system maintains local control and optimization with centralized situational awareness at all times. Such a configuration allows for automated and centrally controlled response.

The system also should be highly scalable, something that can accommodate any device or system that consumes, produces or stores energy. And, it should have the ability to treat each resource as a flexible power and energy asset that can be optimized, aggregated and used for continuous service to the grid. The platform should use real-time information from each DER to calculate optimal setpoints for each resource within the network based on real-time capabilities, local needs and grid needs.

Ideally, the platform also will have software for forecasting as well as energy market interfacing.

To optimize forecasting, the platform should have an adaptive algorithm that learns resource and customer-use characteristics. It also needs to accommodate dynamic conditions, including electricity price, weather, event duration, customer impact, resource constraints and power system status.

And, it should provide forecasting on several time scales, such as 15-minute, hourly, 24-hour and 48-hour intervals.

In addition, the system should be able to combine forecasts of individual resources to create aggregated network forecasts that can be mapped to zones, substations, feeders, transformers and customers. Updates should occur as frequently as every five minutes, and the range of forecasts should be customizable, allowing system operators to see as much as 48 hours ahead of service delivery.

Since we're talking about ancillary services that are often related to energy markets, the aggregation platform should also manage the comprehensive interface to grid operations and market systems. It should be capable of dynamic bidding and evaluation of all available market opportunities. Plus, for recording purposes, it needs to provide real-time measurement and verification, a full audit trail.

With those kinds of network capabilities in place, utilities, ESPs and grid operators can get the same value from a network of DERs that they'd get with a fast-ramping generator or hydro plant. Plus, they get something equally valuable: the opportunity to include all kinds of customers in demand-management and DER-optimization programs.

Such programs allow organizations to cement customer relationships. Rather than merely selling a commodity – electricity – the organization can provide a platform that allows customers to lower peak demand charges and, at the same time, enroll in market programs that actually make the customer some money. Suddenly, the value stack offered the customer includes electricity, revenue earning and cost avoidance.

What's more, the whole thing is far less onerous than traditional load-shedding or demand response initiatives. While DER networks do require customers to drop a certain amount of load at the power-supplier's command, using the right platform, DER participation becomes much less rigid and far more attractive to customers. Done right, there's no discernable interruption to industrial processes or HVAC comfort controls at a facility, because participation can be programmed to happen in accordance to predefined customer constraints.

This approach vastly expands the number of customers who could participate in an intelligent DER-management program, plus it greatly reduces the compliance fatigue when customers are forced to shut down the plant to deliver X-amount of load to the utility per demand response contract.

That, in turn, means that the network can grow more quickly, which feeds the virtuous cycle of Metcalfe's Law: the more, the mightier. It works for telephony, for LinkedIn and for distributed energy resources.