The Atrociously Unfair Data Capacity Tax

And Why It is Time for Capacity-Based Licensing to go Away

What rankles is software licensing that continuously increases in price as the amount of stored data grows.

IT organizations hate licensing software based on capacity. They despise it with a passion for the simple reason that data is growing exponentially, a fact over which they often have little or no control. What rankles is software licensing that continuously increases in price as the amount of stored data grows. Customers are already paying for the actual storage devices to house the data. But capacity-based software metering means they have to pay for it over again and again and again. It has become common for the software licensing to cost more than the actual storage many times over. That doesn't strike most people as a fair bargain. The software is not working any harder, but the price goes up like clockwork every month or year depending on the subscription frequency and data growth. It's maddening to the extreme and IT departments are screaming that it's unsustainable.

Software vendors generally love software licensing metered on capacity. It's simple, easy to track, and very easy to manage. It also increases software revenue without having to do anything, sell anything more, or update the software. That increasing revenue stream makes for a healthy income statement for the vendor. The revenues increase but not their costs. This is why so many software vendors do it. Backup, archiving, replication, data management, data monitoring, and data movement products are generally all addicted to feeding from the capacity metering trough.

There has to be a better way

But here's the rub: IT departments are lashing out. They can't afford what some are calling a legal extortion racket. Since many of these software products are moving data there is a growing belief that a better licensing model should be similar to network bandwidth licensing. Bandwidth is not licensed by the cumulative amount of data that passes through the network. It is licensed by the maximum amount of data that can pass through the network at any given point in time. This is licensing based on performance not capacity. Such licensing is typically based on bits per second. If more data needs to be transported within that same timeframe, more bandwidth is licensed. 100Gbps is licensed at a higher rate than 10Gbps, which is higher than 1Gbps, etc.

More performance is the reason licensing increases. Not the amount of data moved. That makes a lot more sense to the IT departments and it's easier to justify while being fair. The trick is how to do it.

Networking ties software licensing to the hardware on which it is running, or each instance of the software plus the bandwidth. Each switch has its own software license instance. The instance analog for data movement, management, monitoring software can be where the software is running. The performance or bandwidth analog can be tied to the amount of hardware utilized by the software to achieve the desired performance. That can be with physical machines, virtual machines, containers, sockets, or cores. When more performance is required, another software instance can be deployed on additional hardware, or more hardware can be added to the original instances. The problem is, implementing this licensing methodology is more difficult for the software vendor than just doing so on paper. The software itself must be able scale-out to increase performance. The performance scalability should be linear or near linear as more instances on more hardware is introduced. This is a non-trivial effort that will limit the number of software vendors making the conversion to performance-based licensing. But convert they must because the market is demanding it.

Predictable costs over time, with no surprises

This fair and intuitive approach to software licensing is what IT departments worldwide are seeking. We at StrongBox Data Solutions listened and heard their words, their pleadings, and their deeply articulated rejection of capacity-based licensing. As a result, when the StrongLink Autonomous Data Management solution was being developed we included the elimination of capacity-based or other metered licensing as a core design principle in our scale-out architecture. There is no capacity licensing in StrongLink and never will be. It is licensed similarly to networking, where customers purchase the amount of performance they need to match their workflows. Pricing is predictable and unchanging over time for easy budgeting. Much easier than forcing customers to predict how much data they will have under management in five years.

StrongLink software has four simple variables to the StrongLink license structure; 1) Node performance; 2) Number of nodes per site; 3) Whether support for tape is needed; and 4) Whether more than one site is required. For example, each instance of StrongLink running on a customer's server or VM is called a Star or StrongLink node. The Star license amount is based upon the number of CPU cores in that node that are allocated to the system, and which are needed to achieve the desired performance or workflows. For lighter workflows, and simpler use cases, a smaller number of CPU cores would be utilized, and licensed. Heavier or more complex workflows with greater system demands may be allocated a larger number of CPU cores to increase performance of the node.

The second component of the StrongLink license model is whether more than one StrongLink Stars are clustered together for HA or additional performance requirements in a configuration we call a Constellation. With its scale-out architecture StrongLink may be deployed with three or more nodes, with the ability to expand to as many as are needed to match the performance requirements. The third variable is support for tape libraries, which is an optional add-on, since not everyone will need tape. But for those that do, the ability to seamlessly extend their flash and disk-based storage to include on-prem tape archives without capacity taxes provides a huge advantage . The final component is when multiple StrongLink instances in multiple geographic locations are connected remotely, for DR or other replication uses cases. Two or more StrongLink Constellations are called a Galaxy.

Predictable, consistent pricing makes it easy to plan for the future with no surprises, even when it is hard to predict how quickly the data will grow over time. In the same way that a network switch provider would not charge you a variable rate based upon the volume of data going through your pipes, StrongLink licensing helps you get control of your storage costs without adding another "capacity tax" to your environment. This ability to automate policy-based data placement and protection across different storage types, including to tape and cloud, means StrongLink can help you break free of closed vendor ecosystems without being penalized as your data grows.

We invite you to come see for yourself a fairer more cost-effective way to manage both your data and your storage resources. See how StrongLink can help reduce storage and operational costs, so you can get better value from your storage spend, and eliminate the shackles of the "capacity tax".

*To learn more about the StrongLink platform, please visit **www.stronglink.com***

A Solution from STRONGBOX DATA SOLUTIONS

General Customer Contact Info
USA +1 800-643-7148
EU +49 7171 99800-0

info@stronglink.com
www.stronglink.com