Before the

Federal Communications Commission

Washington, DC 20554

In the Matter of	)	
	)	
Amendment of the Commission's Rules with GN Docket No. 12-354 Regard to Commercial Operations in the 3550-3650 MHz Band	)	WT Docket No 12-148
	)	
	)	

1. INTRODUCTION

1.1. About Vanu, Inc.

Vanu, Inc. has it origins in DARPA-sponsored research on software radio initiated at the Massachusetts Institute of Technology. Following the proud tradition of American entrepreneurship, Vanu, Inc. was formed in 1998 to commercialize the research, and has deployed commercial software radio cellular networks in the United States, Canada, India and Nepal. Along the way, Vanu has achieved many industry milestones, including being awarded the first-ever FCC certification for a software defined radio system. Vanu, Inc. sees new creative uses of spectrum, such as those put forth in this NPRM, as essential for speeding innovation and bringing the United States back into prominence in what has become a foreign-dominated wireless infrastructure industry.

2. DISCUSSION

2.1. General Comments

Vanu commends the Commission for putting forth an innovative proposed rulemaking for Commercial Operations in the 3550-3650 MHz Band. By leaving the federal incumbent users in place, this approach holds the potential for freeing up much more spectrum much more rapidly by avoiding the time and cost associated with the conventional spectrum clearing approach. In particular we commend the commission for proposing a three tiered Spectrum Access System (SAS). This approach lets the market decide whether the spectrum will be utilized in a manner similar to a traditional licensed carrier model, or in a way that more closely resembles today's unlicensed operation. While the existing models of dedicated spectrum licensing and unlicensed shared spectrum cannot be directly implemented when there is an incumbent spectrum user, the proposed Tier 2 Priority Access provides many aspects similar to the dedicated model and proposed Tier 3 Generally Authorized Access (GAA) provides something like an unlicensed model. Limiting the SAS to two Tier's would be tantamount to the commission deciding which business model will win in the market place. Furthermore, the proposed three tier model offers the intriguing possibility that the two could co-exist in the same spectrum, bringing the major benefits of each to the market place for the benefit of consumers.

The SAS is the heart of the approach, and many of the concerns of both the federal incumbent users and new entrant commercial users must be dealt with in the design of the SAS. While the system is often referred to as a database or as an extension of the whitespaces database concept, it needs to be a much more active and dynamic system. The system must handle issues related to access control, dynamic

device management and protection of sensitive DoD data. The SAS will also place some requirements on the end user devices, in particular related to certificate management and authentication. While the commission should not specify the design of the SAS, it should specifically call out the major roles and functions that the SAS should serve. The following sections review several key functions of the SAS. Several of the questions raised in the NPRM are addressed in the context of this discussion, since it is ultimately the SAS that is responsible for implementing protocols and automated procedures to address these issues.

2.2. SAS Access Control

Access control is an essential function of the SAS. Devices must contact the SAS for permission to use the spectrum. In order to properly implement access control, a suitable certificate mechanism must be supported by the devices which is discussed in section 2.4. When permission is granted to a device it should have a time-to-live (TTL) associated with it, requiring the device to re-authenticate within the TTL period. This provides a mechanism to effectively revoke permissions from devices found to be operating out of specified parameters and also provides a way of temporarily “clearing spectrum” by not renewing certificates in the case of an event that might invoke priority rights for mission critical users. The appropriate value for the TTL is likely to vary with time and location, and should be set by the SAS in consultation with the federal incumbent users to balance the tension between how quickly malfunctioning devices need to be cleared and not overwhelming the SAS with frequent renewal requests.

Permission for a device to operate may be denied by the SAS for any number of reasons, including: the device is inside of an exclusion zone; the device has been determined to have a flaw and cannot be authorized (or re-authorized) until or fixed by the manufacturer; or there are too many devices in operation in a particular area which could lead to harmful interference to the incumbent. The last point is where Tier 3, GAA, must differ from today’s unlicensed operation in order to protect the incumbent. For example, in the WiFi bands, there was always concern incurring of the “Tragedy of the Commons”, which happens when too many users enter into a given area and raise the interference floor to a level where none of the devices can utilize use the spectrum. In general this has not happened in the WiFi bands primarily because there is a shared interest on everyone’s part to make the spectrum useful. Individuals become their own defacto band managers and make decisions to purchase devices such as phone and a wireless networks hub that operate in different bands. In addition, technology advances such an MIMO have greatly improved the robustness to interference of WiFi routers, allowing many more to coexist in the same space. It is reasonable to assume that future technology advances will further improve the interference tolerance of these systems. This approach does not work when you have an incumbent

spectrum user. Technology advances and individual management techniques may well make the GAA devices able to operate in increasing levels of interference over time, but the incumbent systems, such as radars, are upgraded on a much longer time cycle and will not likely improve their robustness to interference at the same rate. Without access control, more and more GAA devices could be introduced into the band, and they would work quite well, just as most WiFi devices do today. However, the aggregate interference created by the devices could create crippling interference to the incumbent. Instead of Tragedy of the Commons, we could end up with the "Tragedy of the Incumbent", in which case GAA devices function properly but the incumbents would be significantly impaired. The access control mechanism in the SAS must be designed to avoid this outcome. Note that the TTL for GAA users is in conflict with the commissions proposed "term-less" license by rule approach, but we believe it is essential to have a term for the purpose of incumbent interference management, the expectation for renewal for the term GAA use should be very high.

This access control mechanism can also be used to create different zones of operation, rather than just the exclusion zone and one zone of operation. One of the incumbent concerns is that consumer devices might be interfered with or possibly even damaged by some of the high power systems used by the incumbents. The SAS can create multiples zones of operation in which different classes of devices can operate. For example, the default authorization for all devices could be the existing 200 mile exclusion zone. Devices that have been certified for operation with higher tolerances to incumbent signals could then be authorized to operate in zones closer to the incumbent systems. One could imagine a series of concentric circles around the incumbent zone of operation that define different access zones. The zones a device can be authorized for depend upon what certification criteria the device has achieved. The device certification qualifies the robustness to interference of the receiver as well as the survivability of the device under certain types of extreme high power signals. This does not have to be a heavyweight or rigorous certification process; in fact it can and should be a self-certification process, as described in section 2.4. The self certification process is the simplest and imposes the least cost on the device manufacturer. Since the SAS has the ability to revoke permissions for any devices if they are later found out to not meet the stated criteria, this risk incurred through a self certification process is low.

2.3. Dynamic Device Management

It is essential that the SAS be involved in dynamically managing the device permissions. Unlike TV Whitespaces, where the incumbent television stations are stationary and new stations come on and off line on a timescale of years, the spectrum landscape in the proposed band changes on a much more rapid time scale.

The permissions provided by the SAS should specify the frequencies and power levels that the device is allowed to operate at, in addition to area of operation and TTL. The commission has asked what the appropriate power levels are for the various Tiers of operation. With an SAS, this does not have to be a fixed level for all places and all times. For example, in rural areas devices could be allowed to operate at higher power if the device density in the area is low. If more devices are added over time, when the devices operating at the higher power come back to renew their certificate, their allowable power can be lowered. This maximizes spectrum utilization based upon requested usage. Under current static rules, all devices have to abide by the same power level, even if there are relatively few in operation in a given area, which is a waste of spectrum resources. A dynamic approach allows the spectrum usages to be maximized for the number of devices desiring to use the spectrum in a given time period, and to adapt that usage over time to best utilize the spectrum as the number of devices requesting permission changes. If necessary to be consistent with the license-by-rule framework, a very low minimum power (e.g. 100 mW) can be the default for the GAA devices, this limit can then be raised based on coordination with the SAS on a case by case basis.

Dynamic device management can be made far more effective by utilizing sensing data in the device management algorithms. There is concern that pure sensing approaches are not yet viable and they are subject to issues such as the hidden terminal problem. However combing sensing technology that is already utilized in cellular networks today with the SAS will result in a much more effective spectrum management system. Most devices to day perform some sort of spectrum sensing. Cell phones periodically provide measurement reports back to the network containing the identity and strength of the cell towers within range of the phone. Most femto cells have downlink scanning capability built in, and all WiFi access point have scanning capability which is the mechanism they use to find a channel to operate on. Requiring devices in this band to periodically sense the spectrum and report that data back to the spectrum access system will not increase device cost at all, since this capability already today and is needed for effective network and spectrum management even in conventional networks. The sensing data provides timely local information. The SAS contains global information beyond what any particular device is able to sense. The combination can result in a much more spectrum efficient management scheme than either the sensing or SAS can provide individually. Since this kind of reporting of sensing data exists in cellular networks today, this effort can start as a straightforward technology insertion into

the SAS. In the future it can evolve to incorporate newer technologies currently under development, such as those being developed in the DARPA RadioMap program[1].

Dynamic management combined with sensing provides a solution to the boundary problem in the hybrid model in which geographic area licenses would be issued for public property or outdoor areas, while a license-by-rule approach would be employed in private property or indoor areas. The challenge here is in setting an appropriate power level for the license-by-rule users because a single power level does not fit all cases. The appropriate power level for an indoor user in a densely populated apartment building is different from that of a user on a corporate campus or rural area. Dynamic management combined with sensing data provides a solution to this problem. Initially the device can be authorized with a power level based on the other requested use in the geographic area. The default can be a low power limit suitable for indoor urban settings, but if the device is in another region or in an area where very few systems are in use, the initial power limit can be raised. As sensing data comes in from devices in the area, the power level can be adjusted dynamically. For example, if there are Priority Access devices in the area, and they report seeing significant outdoor power levels from a Tier 3 GAA device, the SAS can adjust the power limit of the GAA device downwards. This can be achieved in one of two ways. If the GAA device is network connected, the SAS can immediately send it a message to adjust the allowable power levels. If it is not connected, then the power level is adjusted when the device's TTL expires and the device contacts the SAS for a new certificate. The timeframe required for such adjustments are a factor that should be considered when initially setting the length of the device's TTL. Once the device has been in operation for a while, and the power levels seem suitable, the TTL can be increased when the next certificate is issued to the device.

In summary, we believe that dynamic device management through the SAS is essential to maximize spectrum utilization as the spectrum landscape changes over time, and that it provides an effective mechanism to manage the boundary problem in the hybrid model in which geographic area licenses would be issued for public property or outdoor areas, while a license-by-rule approach would be employed in private property or indoor areas.

[1] DARPA, Strategic Technology Office, Advanced RF Mapping (RadioMap), http://www.darpa.mil/Our_Work/STO/Programs/Advanced_RF_Mapping_(Radio_Map).aspx

2.4. Device Requirements

In order for the SAS to function properly, it must trust that the devices operating in the band will operate within the limits placed on them by the SAS. The trust mechanisms break down into two categories. First, there is the need to authenticate the device and establish a secure communications channel between the device and the SAS. The second is to insure that the device operates with the limits the SAS places upon it, which the SAS may change based upon time and location.

Technical solutions for authentication and establishing a secure channel exist today in systems and standards such as X.509.[2] These are used in a wide range of wireless systems today ranging from WiMAX radios to Android devices and some SDR-based amateur radio systems. The technology for authenticating and securely communicating with a device is mature, cost effective and available today.

The second requirement, to insure the radio operates within the limits specified by the SAS, must be verified as part of the device approval process. The information provided by the SAS includes allowable power levels, frequency ranges, geographic areas and a TTL. The radio itself must have a secure mechanism that does not allow the radio to operate out of bounds in any of these dimensions. This is a small subset of the security issues that have been worked on in SDR research community and in the IEEE standards coordinating committee 41.[3] The SDR problem is much harder, because it involves secure download of software that controls the entire radio operation. In this case, the radio only needs to be designed with a mechanism to limit operation in each of the parameters specified by the SAS, which can be implemented in the radio's firmware and easily tested and verified. We recommend the commission adopt a light weight self-certification process for devices such as self verification or a declaration of conformity so as not to increase time to markets or place a significant cost burden on the device manufacturer. The SAS always has the ability to revoke device operating permission if an issue is discovered, mitigating the need for a more formal certification process involving third parties.

2.5. Protection of sensitive DoD Information

One of the challenges in implementing the SAS is access to sensitive DoD data regarding the location and times of operation of certain systems, and DoD is understandably reluctant to provide all of this data to the multiple spectrum access system vendors. To address this issue, we recommend that the commission

---

[2] Internet X.509 Public Key Infrastructure Certificate and CRL Profile, RFC 2459, http://www.ietf.org/rfc/rfc2459.txt

[3] Formerly IEEE 1900

adopt an approach that utilizes separate databases for sensitive DoD information and public information. The spectrum access systems can then query the DoD database through a secure channel when there are requests for specific locations and frequencies. The DoD database can then respond with the specific information for the relevant access request, without making all of the data accessible all of the time. Furthermore the DoD control allows the DoD to limit access if it suspects the spectrum access systems may have been compromised. We feel this approach is a good balance between DoD's concern at releasing a large amount of sensitive information to a variety of public and private institutions and the need to access specific information to enable spectrum sharing when there are actually systems that need access to the spectrum.

2.6. Indoor Location

One of the challenges with any geographic location-based system is determining position indoors and in other areas where GPS does not function well. The proposed use of geographic area licenses for outdoor areas with a license-by-rule approach employed in private property or indoor areas brings this issue to the forefront. We recommend that the Commission take a "trust but verify" approach to determining location of indoor systems. The "Trust" portion of the approach can be implemented using a method similar to what the FCC has employed for E911 calls over VoIP services can be used.[4] In this approach, the SAS initially relies upon the user or service provider to provide a location where the device is operating. The SAS can then use the sensing data at its disposal to "Verify" this information. The approximate location can be verified using the measurement reports come in to the SAS from the devices in operation. If the reports show signal energy from a device in an area that is far away from the listed location, then the SAS can revoke the device permissions since it can positively determine that the device is not in the location listed. The primary mechanism for this data would be outdoor devices that have GPS capability, so the measurement report contains both a known position and signal strength. Once devices can be verified as being in the listed locations, then their sensing data can also be used to verify other devices in the vicinity.

Respectfully submitted,

/s/ Vanu G. Bose

Vanu G. Bose
CEO and President
Vanu, Inc.

[4] http://www.fcc.gov/guides/voip-and-911-service